arXiv:1603.07069v2 [math.GT] 2 Aug 2016

IMMERSING QUASI-FUCHSIAN SURFACES OF ODD EULER CHARACTERISTIC IN CLOSED HYPERBOLIC 3-MANIFOLDS

YI LIU

Abstract. In this paper, it is shown that every closed hyperbolic 3-manifold contains an immersed quasi-Fuchsian closed subsurface of odd Euler characteristic. The construction adopts the good pants method, and the primary new ingredient is an enhanced version of the connection principle, which allows one to connect any two frames with a path of frames in a prescribed relative homology class of the frame bundle. The existence result is applied to show that every uniform lattice of $\mathrm{PSL}(2,\mathbb{C})$ admits an exhausting nested sequence of sublattices with exponential homological torsion growth. However, the constructed sublattices are not normal in general.

1. Introduction

For any arbitrary closed hyperbolic 3-manifold, immersed quasi-Fuchsian closed subsurfaces have been constructed by J. Kahn and V. Markovic [KM1]. While their subsurfaces are always orientable, it is possible to build non-orientable subsurfaces of even Euler characteristic by modifying their construction, (see [Su1] for example). It remains to be an open question whether a subsurface can be constructed to have odd Euler characteristic, [Ag2, Section 11, Question 7]. The question finds its motivation in several aspects of hyperbolic 3-manifold geometry, including all the results we state in the introduction. The theme result of this paper is an affirmative answer to the existence question:

Theorem 1.1. *Let M be a closed hyperbolic 3-manifold. Then there exists a connected closed surface of odd Euler characteristic Σ which admits a π_1–injective, quasi-Fuchsian immersion $j\colon \Sigma \looparrowright M$.*

When an orientable closed 3-manifold contains an embedded closed subsurface of odd Euler characteristic, the manifold admits a degree-one map onto the real projective 3-space, by a characterization of C. Hayat-Legrand, S. Wang, and H. Zieschang [HWZ, Theorem 4.1]. This fact provides one reason for us to be interested in the virtual existence of such subsurfaces. Since hyperbolic 3-manifold groups are LERF [Ag1, Theorem 9.2], Theorem 1.1 has the following immediate consequence:

Corollary 1.2. *Every closed hyperbolic 3-manifold has an orientable finite cover which admits a degree-one map onto the real projective 3-space.*

Corollary 1.2 might be a toy case of some more general fact. In [Su2], H. Sun has proved that closed hyperbolic 3-manifolds virtually 2-dominates any other orientable closed 3-manifolds. In other words, for any closed orientable 3-manifold N, every closed hyperbolic

Date: August 4, 2016.
2010 *Mathematics Subject Classification.* Primary 57M50; Secondary 57M10, 30F40.
Key words and phrases. good pants, quasi-Fuchsian, homological torsion growth.
Supported by the Recruitment Program of Global Youth Experts of China.

3-manifold M has an orientable finite cover which admits a map onto N of degree 2. It seems that with the improved techniques of this paper (Theorem 3.1), there is a good chance to promote Sun's result to virtual 1-dominations.

Another curious application of Theorem 1.1 shows certain exponential torsion growth for uniform lattices of $\mathrm{PSL}(2,\mathbb{C})$:

Theorem 1.3. *Every uniform lattice* Γ *of* $\mathrm{PSL}(2,\mathbb{C})$ *contains an exhausting nested sequence of non-normal torsion-free sublattices* $\{\Gamma_i\}_{i\in\mathbb{N}}$ *such that*

$$\liminf_{n\to\infty} \frac{\log |H_1(\Gamma_n;\, \mathbb{Z})_{\mathrm{tors}}|}{[\Gamma : \Gamma_n]} > 0.$$

Here the sequence being nested means that each subgroup Γ_n contains its successor Γ_{n+1}, and being exhausting means that the common intersection of all the subgroups is trivial. The term in the logarithmic function is the cardinality of the torsion subgroup of the first integral homology of the group Γ_n.

For general uniform lattices of $\mathrm{PSL}(2,\mathbb{C})$, Theorem 1.3 seems to be the first known exhausting sequence of sublattices with exponential homological torsion growth. A key feature of the construction is that it produces a huge number of disjointly embedded subsurfaces of odd Euler characteristic in the quotient hyperbolic 3-manifolds $\mathbf{H}^3 / \Gamma_n$, so a positive portion of $\mathbb{Z}_2$ homological torsion can be recognized. As to be explained in more details in Section 5, our approach is conceptually simple but technically nontrivial. Besides Theorem 1.1, the construction relies essentially on the virtual specialness of hyperbolic 3-manifold groups, proved by I. Agol [Ag1] and D. Wise [Wi]. In particular, the sublattices are constructed inductively by invoking Wise's Malnormal Special Quotient Theorem. The statement of Theorem 1.3 is known to be false for certain uniform lattices of higher-rank simple real Lie groups, including $\mathrm{SL}(n,\mathbb{R})$ for $n > 2$ and $\mathrm{SO}(p,q)$ for $q > 1$ and large p, [AbGN].

It should be pointed out, however, that Theorem 1.3 does not say much about the well known asymptotic growth conjecture on virtual homological torsion, (see [BV, Conjecture 1.13] for arithmetic lattices and [Lü, Question 13.73] for closed Riemannian manifolds). One (optimistic) version of the conjecture may be stated as follows: Given a lattice Γ of $\mathrm{PSL}(2,\mathbb{C})$, for all exhausting nested sequences of normal torsion-free sublattices of Γ, the limit in Theorem 1.3 equals $\mathrm{Vol}(\mathbf{H}^3 / \Gamma) / 6\pi$. Even if Γ is uniform, the sublattices $\{\Gamma_n\}_{n\in\mathbb{N}}$ of our construction are far from normal. In fact, as n increase, larger and larger balls emerge in the quotient hyperbolic 3-manifolds, but in a highly scattered fashion with relatively small total volume, so the sequence is not even convergent in the Benjamini–Schramm sense.

The major innovation of this paper is an enhanced version of the connection principle in good pants constructions, Theorem 3.1. The reader is referred to Section 2 for terminology and background on this topic. To illustrate the primary issue, suppose that we are asked to construct a good curve γ so that it bounds a surface built up with good pants. It is known that we can design the construction to make γ good and null homologous in M, the oriented closed hyperbolic 3-manifold under consideration. Even so, there is still one geometric obstruction. As a good curve, γ has a so-called canonical lift $\hat{\gamma}$ in the frame bundle $\mathrm{SO}(M)$, which is a loop of frames well defined up to homotopy. For the null homologous good curve γ to bound a good panted subsurface, it turns out that its canonical lift $\hat{\gamma}$ must also be null homologous in $\mathrm{SO}(M)$. Therefore, we need to seek for some refinement of our construction to make sure that the homology class of the canonical lift can be controlled.

A solution of the example task above contains the key idea of our refined construction. An outline is deferred to the last part of the introduction.

It turns out that the construction for Theorem 1.1 can be reduced to a very similar situation: We need to find a 'semi-good' curve $\sqrt{\gamma}$ whose double cover is a bounding good curve γ. In fact, the presence of such a curve γ is necessary since a good panted subsurface of odd Euler characteristic contains an odd number of good pants, so the gluing map as a free involution on the union of cuffs must preserve some component by the parity. One could produce the semi-good curve $\sqrt{\gamma}$ by an *ad hoc* modification of the construction of the example task. Alternatively, as we present in this paper, we can establish a more generally adaptable construction, which allows us to connect frames by good paths of frames in any prescribed relative homology class (Theorem 3.1). Accompanied with the enhanced version of the connection principle, the package of [LM] can be applied to produce good curves and panted 2-complexes in closed hyperbolic 3-manifolds, now with significantly more precise control of their homology. For simplicity, we state the following ordinary form of Theorem 3.1, which roughly says that every integral first homology class of the frame bundle can be represented by a good curve, eventually, as we increase the required length.

Theorem 1.4. *Let M be an oriented closed hyperbolic 3-manifold. Given any homology class $\Xi \in H_1(\mathrm{SO}(M);\mathbb{Z})$ and any constant $\epsilon > 0$, there exists a constant $R_0 = R_0(\epsilon, M, \Xi) > 1$ such that the following holds true. For every constant $R > R_0$, there exists an $(R, \epsilon/R)$–curve γ in M of which the canonical lift $\hat{\gamma}$ represents Ξ in $H_1(\mathrm{SO}(M);\mathbb{Z})$.*

We remark that fixing ϵ and M, the *a priori* bound $R_0(\epsilon, M, \Xi)$ can be chosen to depend linearly on the word length of Ξ, with respect to any finite generating set of the first homology. As R tends to infinity, the number of $(R, \epsilon/R)$–curves representing any given class Ξ grows exponentially fast, in a fashion independent of Ξ. In fact, it is possible to derive the asymptotics from Kimoto–Wakayama [KW], (see also [PS, SW]).

We end up our introduction by sketching a solution to our example task, to construct of a good curve γ with a null homologous canonical lift $\hat{\gamma}$. As mentioned, the first attempt is to create a null homologous good curve ζ. We take two long closed geodesic paths a, b based at the same point $p \in M$, making sure that they point very sharply against each other, with little twist in the normal direction. In other words, we apply the connection principle as usual to construct a and b so that the initial direction of a and the terminal direction of b are approximately some unit vector $\vec{t}_p$ at p, while the terminal direction of a and the initial direction of b are approximately $-\vec{t}_p$; moreover, we require the parallel transport of some unit vector $\vec{n}_p \perp \vec{t}_p$ at p along either a or b is approximately $\vec{n}_p$. Then a null homologous good curve ζ can be constructed as the (reduced) cyclic concatenation of the commutator word $ab\bar{a}\bar{b}$. However, a fundamental calculation tells us that the canonical lift $\hat{\zeta}$ represents the nontrivial element $[\hat{c}] \in H_1(\mathrm{SO}(M);\mathbb{Z})$, in general, (see Lemma 3.4). This should not be too surprising because of the following intuition: If we were on an oriented closed hyperbolic surface rather than in a 3-manifold, the same construction would work perfectly well, but there would be no chance for ζ to bound an immersed subsurface, since it has an odd self-intersection number.

We are therefore hinted to make some essential use of the extra dimension. The trick is to intentionally misalign the invariant normal vectors of a and b with some small angular difference, and take a suitable power of ζ to be the desired γ. For simplicity, let us assume that the injectivity radius of M is not too small, say, at least 1. Choose two unit normal

vectors $\vec{n}_p$ and $\vec{n}_p(2\pi/400)$, the rotation of $\vec{n}_p$ about $\vec{t}_p$ of the angle $2\pi/400$. We construct the closed paths a and b as before except asking a to preserve $\vec{n}_p$ and b to preserve $\vec{n}_p(2\pi/400)$, approximately. This causes some tiny nontrivial honolomy of the concatenated path $ab\bar{a}\bar{b}$: The transport of $\vec{n}_p$ along the path becomes approximately $\vec{n}_p(2\pi/100)$. If we believe, for the intuitive reason above, that the parallel transport path of the frame $\mathbf{p} = (\vec{t}_p, \vec{n}_p, \vec{t}_p \times \vec{n}_p)$ along the commutator path is approximately relatively homologous to the path of spinning frames $\mathbf{p}(\phi)$ about $\vec{t}_p$ parametrized by the angle $\phi \in [0, 2\pi/100]$, then we can expect that the cyclically concatenated path $(ab\bar{a}\bar{b})^{100}$ gives rise to a null homologous good curve γ with a null homologous canonical lift $\hat{\gamma}$, as if the accumulated effect of the normal spinning corrected the homology class of $\hat{\gamma}$ from $[\hat{c}]$ back to 0. To implement the idea, of course, we need to set up fundamental calculations and estimates, and make careful choice of error scales and other constants. The idea condenses roughly to Proposition 3.5, which is the heart of Theorem 3.1.

The paper is organized as follows. In Section 2, we review the good pants construction of Kahn–Markovic and its subsequent development. In Section 3, we introduce our enhanced version of the connection principle. The proof of Theorem 1.4 can be found at the end of that section. Section 4 contains the proof of our theme result, Theorem 1.1. The essential construction for Theorem 1.3 is presented in Section 5, and the complete proof is summarized in Section 6.

Acknowledgements. The author would like to thank Ian Agol and Vlad Markovic for many interesting and inspiring conversations. The author also thanks Hongbin Sun for valuable communications.

2. Preliminaries

This section contains a compact introduction to the good pants method. This method has been invented by J. Kahn and V. Markovic [KM1] to resolve the Surface Subgroup Conjecture and developed by various authors in subsequent works, see [Ha, KM2, LM, Sa, Su1, Su2]. For closed hyperbolic 3-manifolds and potentially for other compact rank-one locally symmetric spaces, it provides a package of tools which enables one to conveniently construct certain π_1–injectively immersed 2-subcomplexes, especially nearly totally geodesic subsurfaces. For the purpose of this paper, we restrict our discussion to closed hyperbolic 3-manifolds and include only a minimal collection of relevant materials.

Let M be an oriented closed hyperbolic 3-manifold. The idea of the good pants method is to produce π_1–injectively immersed 2-complexes in M by gluing pairs of *good pants* in a nice way along their common cuffs, which are *good curves*. Roughly speaking, a good curve in M is an oriented geodesic loop with nearly trivial holonomy; a pair of good pants in M is an oriented immersed pair of pants which is nearly regular and nearly totally geodesic.

In quantitative terms, suppose that $0 < \epsilon < 1$ and $R > 1$ are given constants, where ϵ is presumably very small and R very large. An (R, ϵ)*–curve* is defined to be an oriented geodesic loop of complex length approximately the real value R, with the error required to be at most ϵ in absolute value of the difference. A pair of (R, ϵ)*–pants* is defined to be an (unmarked) oriented immersed pair of pants with the following extra requirements on its shape: The three cuffs are required to be (R, ϵ)–curves; the three seams are required to be the unique arcs of the pair of pants which connect between the cuffs as common perpendicular geodesic arcs; as hence there are six well-defined feet, namely, the unit normal vectors of the

cuffs attached at the endpoints of the seams pointing into the seams, the final requirement is that each foot of a cuff should be approximately equal to the parallel transport of the other foot of that cuff, along the half cuff in the forward direction, and that the error should be at most ϵ measured in angle between unit normal vectors. In other words, it only makes sense to speak of the complex *half* length of an (R,ϵ)–curve as a cuff of a (nonsingular) π_1–injectively immersed pair of pants, and being a pair of good pants requires all the cuffs to have complex half length ϵ–close to $R/2$ (rather than $R/2+\pi\cdot\sqrt{-1}$). We do not distinguish (R,ϵ)–curves which are the same up to change of parametrization, or (R,ϵ)-pants which are the same up to homotopy and orientation-preserving self-homeomorphism. The reader is referred to [KM1, Section 2] for the original definition and [LM, Section 2] for expanded discussions.

Adopt the notations

$$\mathbf{\Gamma}_{R,\epsilon}(M) = \{\,(R,\epsilon)\text{–curves of } M\};\ \mathbf{\Pi}_{R,\epsilon}(M) = \{\,(R,\epsilon)\text{–pants of } M\,\}.$$

For any given constant ϵ, as long as R is sufficiently large, where an *a priori* lower bound depends only on M and ϵ, $\mathbf{\Gamma}_{R,\epsilon}(M)$ and $\mathbf{\Pi}_{R,\epsilon}(M)$ are always non-empty. In fact, (R,ϵ)–curves and (R,ϵ)–pants can be produced using the following basic construction, which is an immediate consequence of the (exponential) mixing property of the frame flow, (see [KM1, Lemma 4.4], [LM, Lemma 4.15]):

Proposition 2.1 (Connection principle)**.** *Let M be an oriented closed hyperbolic 3-manifold. Let $\vec{t}_p \perp \vec{n}_p$ and $\vec{t}_q \perp \vec{n}_q$ be two pairs of orthogonal unit vectors at the points p and q of M, respectively.*

Given any positive constant δ, and for every sufficiently large positive constant L with respect to M and δ, there exists a geodesic path s in M from p to q with the following properties:

- *The length of s is δ–close to L. The initial direction of s is δ–close to $\vec{t}_p$ and the terminal direction of s is δ–close to $\vec{t}_q$, where the distance is measured by the angle between unit vectors.*
- *The parallel transport from p to q along s takes $\vec{n}_p$ to a unit vector $\vec{n}'_q$ which is δ–close to $\vec{n}_q$.*

Moreover, as the required cuff length R grows, there will eventually be plenty of (R,ϵ)–pants and their feet tend to be very evenly distributed on the unit normal bundle over every (R,ϵ)–curve. By nicely gluing a suitable collection of good pants along common cuffs, Kahn and Markovic are able to resolve the Surface Subgroup Conjecture for closed hyperbolic 3-manifolds [KM1]:

Theorem 2.2. *Every closed hyperbolic 3-manifold contains a π_1–injectively immersed quasi-Fuchsian closed subsurface.*

Associated with the constructed surface is a naturally induced pants decomposition with markings, so the shape of surface can be described by its complex Fenchel–Nielsen coordinates. In these terms, the complex half length and the shearing parameter for each glued cuff C are approximately $(\mathbf{hl}(C), s(C)) \approx (R/2, 1)$. The first component has error at most ϵ and the second component has error at most ϵ/R.

The relative version of the construction problem has been studied in [LM], namely, whether a collection of good curves can bound a subsurface which is nicely glued from good pants. Regarding any given collection of (R,ϵ)–curves (possibly with multiplicity) as an element of

the free integral module $\mathbb{Z}\mathbf{\Gamma}_{R,\epsilon}(M)$ (in the non-negative hyper-octant), and similarly (R,ϵ)–pants as of $\mathbb{Z}\mathbf{\Pi}_{R,\epsilon}(M)$, an obstruction to solving the relative problem lies in the cokernel of the homomorphism:

$$\partial\colon \mathbb{Z}\mathbf{\Pi}_{R,\epsilon}(M) \longrightarrow \mathbb{Z}\mathbf{\Gamma}_{R,\epsilon}(M),$$

defined by taking any pair of pants to the sum of its three cuffs. For large R with respect to ϵ and M, the (R,ϵ)-*panted cobordism group* of M is defined to be:

$$\mathbf{\Omega}_{R,\epsilon}(M) \;=\; \mathbb{Z}\mathbf{\Gamma}_{R,\epsilon}(M) \,/\, \partial(\mathbb{Z}\mathbf{\Pi}_{R,\epsilon}(M)).$$

The group $\mathbf{\Omega}_{R,\epsilon}(M)$ can be fully characterized by the following correspondence [LM, Theorem 5.2]. We denote by $\mathrm{SO}(M)$ the $\mathrm{SO}(3)$–principal bundle over M of orthonormal frames with right orientation.

Theorem 2.3. *For any sufficiently small positive constant ϵ with respect to M, and any sufficiently large positive constant R with respect to M and ϵ, there exists a canonical isomorphism:*

$$\Phi\colon \mathbf{\Omega}_{R,\epsilon}(M) \longrightarrow H_1(\mathrm{SO}(M);\,\mathbb{Z}).$$

Moreover, for all $\gamma \in \mathbf{\Gamma}_{R,\epsilon}(M)$, the projection of $\Phi(\gamma)$ in $H_1(M;\,\mathbb{Z})$ equals the homology class of γ.

For the goal of this paper, it is important to understand the implementation of Φ. It suffices to specify the assignment of Φ to each curve $\gamma \in \mathbf{\Gamma}_{R,\epsilon}(M)$. In fact, the homology class $\Phi(\gamma)$ can be represented by the *canonical lift* of the curve γ, denoted as

$$\hat{\gamma} \in \pi_1(\mathrm{SO}(M)),$$

and $\hat{\gamma}$ is defined as follows.

Suppose that γ is an (R,ϵ)–curve. Take a special orthonormal frame $\mathbf{p} = (\vec{t}_p, \vec{n}_p, \vec{t}_p \times \vec{n}_p) \in \mathrm{SO}(M)$ at a point $p \in \gamma$ such that $\vec{t}_p$ is the unit tangent vector of γ at p. Here the cross notation stands for the cross product for the oriented tangent space T_pM with the Riemannian inner product. For sufficiently large R with respect to M and ϵ, the parallel transport of $\mathbf{p}$ along γ back to p is a frame $\mathbf{p}'$ which can be connected with $\mathbf{p}$ by a unique shortest path in $\mathrm{SO}(M)|_p$. Then the canonical lift $\hat{\gamma}$ can be represented by the closed path which is the concatenation of three consecutive subpaths: The first is the parallel-transportation path from $\mathbf{p}$ to $\mathbf{p}'$ along γ; the second is a closed path in $\mathrm{SO}(M)|_p$ based at $\mathbf{p}'$ which represents the unique nontrivial central element

$$\hat{c} \in \pi_1(\mathrm{SO}(M));$$

and the third is the shortest path in $\mathrm{SO}(M)|_p$ connecting $\mathbf{p}'$ and $\mathbf{p}$. The resulting loop of frames does not depend on the auxiliary choices up to free homotopy in $\mathrm{SO}(M)$, so $\hat{\gamma} \in \pi_1(\mathrm{SO}(M))$ is well defined. Note that the second sub-path above ensures that Φ vanishes on boundary of pants.

It turns out that elements of $\mathbf{\Omega}_{R,\epsilon}(M)$ are the only obstructions to solving the relative construction problem. For simplicity, we state a special case as follows. The case with multicurve boundary is completely analogous, see [LM], [Su2, Corollary 2.7].

Theorem 2.4. *For any sufficiently small positive constant ϵ with respect to M, and any sufficiently large positive constant R with respect to M and ϵ, the following statement holds true.*

For any curve $\gamma \in \mathbf{\Gamma}_{R,\epsilon}(M)$, there exists an oriented, connected, π_1–injectively immersed quasi–Fuchsian subsurface of M which is (R,ϵ)–panted and bounded by γ if and only if the canonical lift $\hat{\gamma} \in \pi_1(M)$ is null homologous in $\mathrm{SO}(M)$.

Note that Theorem 2.4 does not say anything about the existence of (R,ϵ)-curves with homologically trivial canonical lifts. The existence follows from Theorem 1.4, which is proved in the rest of this paper.

We finish this section with the following finer version of the good pants method that we have discussed so far. This has been pointed out by [Su2, Remark 2.9], see also [Sa], as a consequence of the exponential mixing rate of the frame flow. In the rest of this paper, we adopt the finer version so as to invoke some estimations from [Su1].

Declaration 2.5. All the notations and results that have appeared in this section remain compatible and true with the constant ϵ replaced by ϵ/R, (and δ replaced by δ/L).

3. Connecting frames in a prescribed homology class

Let M be an oriented closed hyperbolic 3-manifold. Denote by $\mathrm{SO}(M)$ the special orthonormal frame bundle over M. Let $\mathbf{p}$ and $\mathbf{q}$ be a pair of (possibly coincident) frames in $\mathrm{SO}(M)$ at points p and q of M, respectively. Denote by

$$\pi_1(\mathrm{SO}(M), \mathbf{p}, \mathbf{q})$$

the set of relative homotopy classes of paths in $\mathrm{SO}(M)$ from $\mathbf{p}$ to $\mathbf{q}$. Each element $\hat{\xi}$ of $\pi_1(\mathrm{SO}(M), \mathbf{p}, \mathbf{q})$ represents a relative homology class $[\hat{\xi}] \in H_1(\mathrm{SO}(M), \mathbf{p} \cup \mathbf{q}; \mathbb{Z})$ with the property $\partial_*[\hat{\xi}] = [\mathbf{q}] - [\mathbf{p}]$, where

$$\partial_* \colon H_1(\mathrm{SO}(M), \mathbf{p} \cup \mathbf{q}; \mathbb{Z}) \longrightarrow H_0(\mathbf{p} \cup \mathbf{q}; \mathbb{Z})$$

is the induced homomorphism of the relative pair $(\mathrm{SO}(M), \mathbf{p} \cup \mathbf{q})$. Since the preimage $\partial_*^{-1}([\mathbf{q}] - [\mathbf{p}])$ is an affine $H_1(\mathrm{SO}(M); \mathbb{Z})$, one may think of the assignment $\hat{\xi} \mapsto [\hat{\xi}]$ as a model of a natural relative grading for $\pi_1(\mathrm{SO}(M), \mathbf{p}, \mathbf{q})$ such that the grading differences between elements are valued in $H_1(\mathrm{SO}(M); \mathbb{Z})$. In this way, $\pi_1(\mathrm{SO}(M), \mathbf{p}, \mathbf{q})$ decomposes naturally into the disjoint union of its grading classes.

The following theorem is an enhanced version of the connection principle which respects gradings.

Theorem 3.1. *Let M be an oriented closed hyperbolic 3-manifold. Let $\mathbf{p} = (\vec{t}_p, \vec{n}_p, \vec{t}_p \times \vec{n}_p)$ and $\mathbf{q} = (\vec{t}_q, \vec{n}_q, \vec{t}_q \times \vec{n}_q)$ in $\mathrm{SO}(M)$ be a pair of special orthonormal frames at points p and q of M, respectively. Let $\Xi \in H_1(\mathrm{SO}(M), \mathbf{p} \cup \mathbf{q}; \mathbb{Z})$ be a relative homology class with boundary $\partial_*[\hat{\xi}] = [\mathbf{q}] - [\mathbf{p}]$.*

Given any positive constant δ, and for every sufficiently large positive constant L with respect to M, Ξ, and δ, there exists a geodesic path s in M from p to q with the following properties:

- *The length of s is (δ/L)–close to L. The initial direction of s is (δ/L)–close to $\vec{t}_p$ and the terminal direction of s is (δ/L)–close to $\vec{t}_q$.*
- *The parallel transport from p to q along s takes $\mathbf{p}$ to a frame $\mathbf{q}'$ which is (δ/L)–close to $\mathbf{q}$, and there exists a unique shortest path in $\mathrm{SO}(M)|_q$ between $\mathbf{q}'$ and $\mathbf{q}$.*
- *Denote by $\hat{s}$ the path which is the concatenation of the parallel-transport path from $\mathbf{p}$ to $\mathbf{q}'$ with the shortest path from $\mathbf{q}'$ to $\mathbf{q}$. The relative homology class represented by $\hat{s} \in \pi_1(\mathrm{SO}(M), \mathbf{p}, \mathbf{q})$ equals Ξ.*

Note that when $\mathbf{p}$ coincides with $\mathbf{q}$, the union $\mathbf{p} \cup \mathbf{q}$ is understood as a single point in $\mathrm{SO}(M)$.

In the rest of this section, we prove Theorem 3.1 by construction. We fix an oriented closed hyperbolic 3-manifold M throughout this section.

3.1. **Basic calculations for parallel-transport paths of frames.** We start by three calculations for the homology classes of frame paths. The first calculation tells us that if we concatenate a consecutive chain of long geodesic paths in M with small total bending, then parallel transport of any frame along the concatenated path or its reduction yields almost the same path of frames up to homotopy. The second calculation is parallel to the first one, for the reduction of the cyclic concatenation of a consecutive cycle. The third calculation considers a special situation that if we concatenate a consecutive commutator chain of long geodesic paths in M with small bending and twisting, then parallel transport of any frame along the concatenated path yields a path of frames which is almost closed and null homologous in the frame bundle.

By a *consecutive chain* of geodesic paths in M, we mean a finite sequence of oriented geodesic paths such that the terminal endpoint of any member of the sequence, except the final one, is the initial endpoint of its successor. The *reduced concatenation* of a consecutive chain is the unique geodesic path in M which is homotopic to the concatenation of the chain relative to the endpoints. A *consecutive cycle* of geodesic paths is a chain whose last member has its terminal endpoint the same as the initial point of the first one. The *reduced cyclic concatenation* is the unique geodesic loop without base point which is freely homotopic to the cyclic concatenation of the cycle.

We do not attempt to make the most economic choices for universal constants in the estimates. Instead, powers of ten are often used, and the power roughly counts the number of basic steps that have been taken.

Lemma 3.2. *Let $s_1, \cdots, s_m$ be a consecutive chain of m geodesic paths in M. Suppose that the length of each s_i is at least L, and suppose that the terminal direction of each s_i, except s_m, is δ–close to the initial direction of s_{i+1}. Denote by s the reduced concatenation of $s_1, \cdots, s_m$. Let $\mathbf{p}$ be a frame in $\mathrm{SO}(M)$ at the initial endpoint p of s_1. Denote by $\mathbf{q}$ and $\mathbf{q}'$ the parallel transport of $\mathbf{p}$ along the concatenation of $s_1, \cdots, s_m$ and along s to the terminal endpoint q of s_m, respectively.*

Then for universally small $(m\delta)$ and sufficiently large L with respect to δ, there exists a unique shortest path in $\mathrm{SO}(M)|_q$ from $\mathbf{q}'$ to $\mathbf{q}$. Moreover, the parallel-transport path from $\mathbf{p}$ to $\mathbf{q}$ is homotopic to the concatenation of the parallel-transport path from $\mathbf{p}$ to $\mathbf{q}'$ and the shortest path in $\mathrm{SO}(M)|_q$ from $\mathbf{q}'$ to $\mathbf{q}$.

Proof. First consider the basic case if m equals 2. Then the consecutive geodesic paths s_1, s_2 spans a unique geodesic 2-simplex T, which is immersed in M, and the reduced concatenated path s is geometrically the third edge of T. For universally small δ and sufficiently large L with respect to δ, the area of T can be bounded by 2δ. As parallel transport takes any frame $\mathbf{p} \in \mathrm{SO}(M)|_p$ along the chain s_1, s_2 to $\mathbf{q} \in \mathrm{SO}(M)|_q$, and along s to $\mathbf{q}' \in \mathrm{SO}(M)|_q$, it follows that the distance between $\mathbf{q}$ and $\mathbf{q}'$ can be bounded by 10δ. We may in addition require that δ is so small that that any ball of radius 100δ in $\mathrm{SO}(M)|_q$ is embedded. Then there is a unique path of the shortest length in $\mathrm{SO}(M)|_q$ connecting $\mathbf{q}'$ and $\mathbf{q}$, which we denote as $\hat{\eta}$.

To see the claimed homotopy in this case, let $h\colon [0,1] \to T$ be the altitude of T on the side s so that $h(1)$ is the joint point of s_1 and s_2, and $h(0)$ lies on s. For a parameter $t \in [0,1]$,

consider the t-family of consecutive chains $s_1(h(t)), s_2(h(t))$ in T which is formed by the two geodesic segments $[p, h(t)]$ and $[h(t), q]$. Denote by $\hat{\xi}_t$ the path of frames given by the parallel transport of $\mathbf{p}$ along the concatenation of $s_1(h(t)), s_2(h(t))$. As $t \in [0, 1]$ varies, the endpoint of $\hat{\xi}_t$ in $\mathrm{SO}(M)|_q$ gives rise to a path $\hat{\eta}'$ from $\mathbf{q}'$ to $\mathbf{q}$. It follows from the construction of $\hat{\xi}_t$ that $\hat{\xi}_0\hat{\eta}'$ is homotopic to $\hat{\xi}_1$, the parallel-transport path along the concatenated chain s_1, s_2. On the other hand, the distance estimation as above can also be applied for pairs of points on $\hat{\eta}'$, so $\hat{\eta}$ has diameter bounded by 10δ in $\mathrm{SO}(M)|_q$. It follows that $\hat{\eta}'$ can be homotoped to $\hat{\eta}$ within $\mathrm{SO}(M)|_q$. Then $\hat{\xi}_1$ is homotopic to $\hat{\xi}_0\hat{\eta}$, which proves the basic case.

The general case can be done by applying the basic case for $m-1$ times. In other words, we triangulate the polygon formed by the consecutive chain $s_1, \cdots, s_m$ and the reduced concatenation s by adding diagonals, and then apply the basic case to each geodesic 2-simplex that fills a triangle. (See [LM, Lemma 4.8 (a)] for the geometry of the diagonals.) Then the total area of the 2-simplices is bounded by about $m\delta$. When this quantity is small enough, $\mathbf{q}'$ and $\mathbf{q}$ can be connected by a uniquely shortest path in $\mathrm{SO}(M)|_q$, of length at most $100m\delta$. We omit the details because the estimation is straightforward. □

Lemma 3.3. *Let $s_1, \cdots, s_m$ be a consecutive cycle of m geodesic paths in M. Suppose that the length of each s_i is at least L, and suppose that the terminal direction of each s_i is δ–close to the initial direction of s_{i+1}, where s_{m+1} means s_1. Denote by γ the reduced cyclic concatenation of $s_1, \cdots, s_m$. Let $\mathbf{p}$ be a frame in $\mathrm{SO}(M)$ at the initial endpoint p of s_1. Denote by $\mathbf{p}'$ the parallel transport of $\mathbf{p}$ along the cyclic concatenation of $s_1, \cdots, s_m$ back to p. Let $\mathbf{q}$ be a frame in $\mathrm{SO}(M)$ at a point $q' \in \gamma$. Denote by $\mathbf{q}'$ the parallel transport of $\mathbf{q}$ along γ back to q.*

Then for universally small $(m\delta)$ and sufficiently large L with respect to δ, there exist unique shortest paths in $\mathrm{SO}(M)|_p$ from $\mathbf{p}'$ to $\mathbf{p}$, and in $\mathrm{SO}(M)|_q$ from $\mathbf{q}'$ to $\mathbf{q}$. Moreover, the cyclic concatenation of parallel-transport path from $\mathbf{p}$ to $\mathbf{p}'$ with the shortest path in $\mathrm{SO}(M)|_p$ from $\mathbf{p}'$ to $\mathbf{p}$ is freely homotopic to the cyclic concatenation of the parallel-transport path from $\mathbf{q}$ to $\mathbf{q}'$ with and the shortest path in $\mathrm{SO}(M)|_q$ from $\mathbf{q}'$ to $\mathbf{q}$.

Proof. Note that under the assumptions, the annulus between the cyclic concatenation of the cycle and its reduction can be chosen to have small area, and $q \in \gamma$ can be chosen 10δ–close to p. See [LM, Lemma 4.8 (b)] for the geometry. Then the lemma can be proved in a way similar to Lemma 3.2. □

Lemma 3.4. *Let a, b be a pair of geodesic paths in M with all the endpoints the same point p. Suppose that there exists a unit vector $\vec{t}$ at $\vec{p}$ such that $\vec{t}$ is δ–close to the initial direction of a and the terminal direction of b, and that $-\vec{t}$ is δ–close to the terminal direction of a and the initial direction of b. Moreover, suppose that there exists a unit vector $\vec{n} \perp \vec{t}$ at p of which the parallel transport along a and along b are both δ–close to $\vec{n}$. Let $\mathbf{p}$ be a frame in $\mathrm{SO}(M)$ at p. Denote by $\mathbf{q}$ the parallel transport of $\mathbf{p}$ along the concatenated path $ab\bar{a}\bar{b}$, where the bar notation stands for the orientation reversal.*

Then for universally small δ, there exists a unique shortest path in $\mathrm{SO}(M)|_p$ from $\mathbf{q}$ to $\mathbf{p}$. Moreover, the concatenation of the parallel-transport path from $\mathbf{p}$ to $\mathbf{q}$ and the shortest path in $\mathrm{SO}(M)|_p$ from $\mathbf{q}$ to $\mathbf{p}$ is null homologous in $\mathrm{SO}(M)$.

Proof. Observe that it suffices to prove the lemma for a specific frame $\mathbf{p}$ in $\mathrm{SO}(M)|_p$. In fact, suppose that we have found a (cellular) 2-chain $C = k_1\sigma_1 + \cdots + k_r\sigma_r$ bounded by the claimed cycle for $\mathbf{p}$, where $\sigma_i \colon D^2 \to \mathrm{SO}(M)$ are 2-cells. Since any other frame $\mathbf{p}'$ in $\mathrm{SO}(M)|_p$ can be

written as $\mathbf{p}\cdot g$ where $g \in \mathrm{SO}(3)$ is a constant matrix, the 2-chain $C\cdot g = k_1\sigma_1\cdot g+\cdots+k_r\sigma_r\cdot g$ is bounded by the claimed cycle for $\mathbf{p}'$. This is because the right action of $\mathrm{SO}(3)$ on the principal bundle $\mathrm{SO}(M)$ preserves the metric and commutes with parallel transport.

In the rest of the proof, we argue for the frame $\mathbf{p} = (\vec{t}, \vec{n}, \vec{t}\times\vec{n})$ in $\mathrm{SO}(M)|_p$. Denote by $\mathbf{p}^\dagger \in \mathrm{SO}(M)$ the frame $(-\vec{t}, \vec{n}, -\vec{t}\times\vec{n})$.

For universally small δ, we may assume that any ball of radius 1000δ in $\mathrm{SO}(M)|_p$ is embedded. By the assumption, the parallel transport of $\mathbf{p}$ consequentially along a, b, $\bar{a}$, $\bar{b}$ gives rise to four frames $\mathbf{p}_1, \cdots, \mathbf{p}_4 \in \mathrm{SO}(M)|_p$ such that $\mathbf{p}_1$ and $\mathbf{p}_3$ lie in the (100δ)–neighborhood of $\mathbf{p}^\dagger$, while $\mathbf{p}_2$ and $\mathbf{p}_4$ lie in the (100δ)–neighborhood of $\mathbf{p}$. Note that $\mathbf{p}_4 = \mathbf{q}$, and it is already clear that $\mathbf{q}$ can be connected by a unique shortest path in $\mathrm{SO}(M)|_p$.

It remains to show that the claimed cycle is a boundary in $\mathrm{SO}(M)$. Writing $\mathbf{p}_0 = \mathbf{p}$, the claimed cycle as the sum of the parallel-transport paths $[\mathbf{p}_i, \mathbf{p}_{i+1}]$ for $i = 0, 1, 2, 3$ and the shortest path $[\mathbf{p}_4, \mathbf{p}_0]$ in $\mathrm{SO}(M)|_p$. We argue by homologically simplifying the claimed cycle until it is obviously null homologous in $\mathrm{SO}(M)$. To this end, denote by

$$\mathbf{p}\cdot R_N(\theta) \in \mathrm{SO}(M)|_p$$

the rotation of $\mathbf{p}$ about $\vec{n}$ by an angle θ, which is given by the right action of the matrix

$$R_N(\theta) \;=\; \begin{bmatrix} \cos\theta & 0 & \sin\theta \\ 0 & 1 & 0 \\ -\sin\theta & 0 & \cos\theta \end{bmatrix} \in \mathrm{SO}(3).$$

Note that $\mathbf{p}\cdot R_N(0) = \mathbf{p}$ and $\mathbf{p}\cdot R_N(\pi) = \mathbf{p}^\dagger$. Denote by $\mathbf{p}_i^\dagger = \mathbf{p}_i\cdot R_N(\pi)$, and by $\hat{\xi}_i$ the path from $\mathbf{p}_i$ to $\mathbf{p}_i^\dagger$, parametrized as $\mathbf{p}\cdot R_N(\theta)$ for $\theta \in [0, \pi]$. We have approximately

$$\mathbf{p}_0, \mathbf{p}_1^\dagger, \mathbf{p}_2, \mathbf{p}_3^\dagger, \mathbf{p}_4 \approx \mathbf{p}$$

with error at most 100δ from $\mathbf{p}$, and similarly,

$$\mathbf{p}_0^\dagger, \mathbf{p}_1, \mathbf{p}_2^\dagger, \mathbf{p}_3, \mathbf{p}_4^\dagger \approx \mathbf{p}^\dagger.$$

By our assumption on the smallness of δ, those frames near $\mathbf{p}$, or $\mathbf{p}^\dagger$ respectively, can be mutually connected by unique shortest paths in $\mathrm{SO}(M)|_p$.

Consider the paths $[\mathbf{p}_0, \mathbf{p}_1]$ and $[\mathbf{p}_2, \mathbf{p}_3]$. Because there is a rectangle parametrized as a family of parallel-transport paths $[\mathbf{p}_2, \mathbf{p}_3]\cdot R_N(\theta)$ along $\bar{a}$ where $\theta \in [0, \pi]$, the path $[\mathbf{p}_2, \mathbf{p}_3]$ is homologous to the (1-chain) sum of $\hat{\xi}_2$ and $[\mathbf{p}_2, \mathbf{p}_3]\cdot R_N(\pi)$ and $-\hat{\xi}_3$. The middle term $[\mathbf{p}_2, \mathbf{p}_3]\cdot R_N(\pi)$ equals the parallel-transport path from $\mathbf{p}_2^\dagger$ along $\bar{a}$, where $\mathbf{p}_2^\dagger$ is approximately $\mathbf{p}^\dagger$. Since $[\mathbf{p}_0, \mathbf{p}_1]$ is the parallel-transport path from $\mathbf{p}_0$ along a, where $\mathbf{p}_0$ is approximately $\mathbf{p}$, the path $[\mathbf{p}_2, \mathbf{p}_3]\cdot R_N(\pi)$ is almost the orientation reversal of $[\mathbf{p}_0, \mathbf{p}_1]$. More precisely, their sum is homologous to the sum of the shortest paths $[\mathbf{p}_0, \mathbf{p}_2^\dagger]$ and $[\mathbf{p}_3^\dagger, \mathbf{p}_1]$ near $\mathbf{p}$ and $\mathbf{p}^\dagger$, respectively. This yields

$$[\mathbf{p}_0, \mathbf{p}_1] + [\mathbf{p}_2, \mathbf{p}_3] = \hat{\xi}_2 - \hat{\xi}_3 + [\mathbf{p}_0, \mathbf{p}_2^\dagger] + [\mathbf{p}_3^\dagger, \mathbf{p}_1],$$

as 1-chains of $\mathrm{SO}(M)$ modulo 1-boundaries. Similarly,

$$[\mathbf{p}_1, \mathbf{p}_2] + [\mathbf{p}_3, \mathbf{p}_4] = \hat{\xi}_3 - \hat{\xi}_4 + [\mathbf{p}_1, \mathbf{p}_3^\dagger] + [\mathbf{p}_4^\dagger, \mathbf{p}_2],$$

modulo 1-boundaries. Since $\mathbf{p}_2$ and $\mathbf{p}_4$ are near to each other,

$$\hat{\xi}_2 - \hat{\xi}_4 = -[\mathbf{p}_2^\dagger, \mathbf{p}_4^\dagger] - [\mathbf{p}_4, \mathbf{p}_2],$$

modulo 1-boundaries. Therefore, the claimed cycle

$$[\mathbf{p}_0, \mathbf{p}_1] + [\mathbf{p}_1, \mathbf{p}_2] + [\mathbf{p}_2, \mathbf{p}_3] + [\mathbf{p}_3, \mathbf{p}_4] + [\mathbf{p}_4, \mathbf{p}_0]$$

is equal to a linear combination of shortest paths in $\mathrm{SO}(M)|_p$ that are near $\mathbf{p}$ or $\mathbf{p}^\dagger$, modulo 1-boundaries of $\mathrm{SO}(M)$. It follows that the claimed cycle is null homologous in $\mathrm{SO}(M)$. □

3.2. **Substitution for paths of spinning frames.** Given a parameter $\phi \in \mathbb{R}$, and for any frame $\mathbf{p} = (\vec{t}_p, \vec{n}_p, \vec{t}_p \times \vec{n}_p)$ in $\mathrm{SO}(M)$ at a point p in M, we denote by

$$\mathbf{p}(\phi) \;=\; (\vec{t}_p, \vec{n}_p(\phi), \vec{t}_p \times \vec{n}_p(\phi)) \in \mathrm{SO}(M)|_p$$

the rotation of $\mathbf{p}$ about $\vec{t}_p$ by an angle ϕ. In other words,

$$\mathbf{p}(\phi) \;=\; \mathbf{p} \cdot R_T(\phi),$$

where $R_T(\phi) \in \mathrm{SO}(3)$ stands for the matrix

$$R_T(\phi) \;=\; \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos\phi & -\sin\phi \\ 0 & \sin\phi & \cos\phi \end{bmatrix}.$$

This gives rise to a 1-parameter family of spinning frames

$$\mathbf{p} \cdot R_T \colon \mathbb{R} \to \mathrm{SO}(M)|_p.$$

Proposition 3.5. *Let $\mathbf{p} = (\vec{t}_p, \vec{n}_p, \vec{t}_p \times \vec{n}_p)$ in $\mathrm{SO}(M)$ be a frame at a point p in M. Let $\phi \in \mathbb{R}$ be a constant and denote by $\hat{\omega}$ the path of spinning frames from $\mathbf{p}$ to $\mathbf{p}(\phi)$ which is parametrized by the interval $[0, \phi]$.*

Given any positive constant δ, and for every sufficiently large L with respect to M and δ, there exists a geodesic path s in M with both endpoints p which satisfies the following requirements:

- *The length of s is (δ/L)–close to L. The initial direction and the terminal direction of s are both (δ/L)–close to $\vec{t}_p$.*
- *The parallel transport from p back to p along s takes $\mathbf{p}$ to a frame $\mathbf{q}'$ which is (δ/L)–close to $\mathbf{p}(\phi)$, and there exists a unique shortest path in $\mathrm{SO}(M)|_p$ between $\mathbf{q}'$ and $\mathbf{p}(\phi)$.*
- *Denote by $\hat{s}$ the path which is the concatenation the parallel-transport path from $\mathbf{p}$ to $\mathbf{q}'$ with the shortest path from $\mathbf{q}'$ to $\mathbf{p}(\phi)$. The relative homology class represented by $\hat{s} \in \pi_1(\mathrm{SO}(M), \mathbf{p}, \mathbf{p}(\phi))$ equals $[\hat{\omega}] \in H_1(\mathrm{SO}(M), \mathbf{p} \cup \mathbf{p}(\phi); \mathbb{Z})$.*

3.2.1. *The local case.* Supposing that ϕ is given with $|\phi|$ small enough, we first show a local case of Proposition 3.5, namely, for any $\delta > 0$ at most $|\phi|$, and any for sufficiently large $L > 0$ with respect to M and δ, the geodesic path s in M can be constructed with the asserted properties of Proposition 3.5.

To be specific, it suffices to require $|\phi|$ to be so small that the $(|\phi| \times 10^4)$–neighborhood of $\mathbf{p}$ is embedded in $\mathrm{SO}(M)|_p$. Then the path of spinning frames $\mathbf{p} \cdot R_T \colon [0, 10\phi] \to \mathrm{SO}(M)|_p$ is the unique shortest path between $\mathbf{p}$ and $\mathbf{p}(10\phi)$ in $\mathrm{SO}(M)|_p$. Then for any $0 < \delta < |\phi|$, we require $L > 1$ to be so large that the connection principle can be applied with respect to $\delta \times 10^{-3}$.

The construction of the local case is as follows. By the connection principle, we construct geodesic paths a and b in M with all endpoints p of the following shape (or pose):

- The length of a is $(\delta/100L)$–close to $L/4$. The same holds for the length of b.

- The initial direction and the terminal direction of a are $(\delta/100L)$–close to $\vec{t}_p$ and $-\vec{t}_p$ respectively. Oppositely, the initial direction and the terminal direction of b are $(\delta/100L)$–close to $-\vec{t}_p$ and $\vec{t}_p$ respectively.
- The parallel transport along a takes $\vec{n}_p$ back to $\vec{n}_p$ up to an error at most $\delta/100L$ in angle. However, the parallel transport along b takes $\vec{n}_p$ to $\vec{n}_p(\phi/2)$ with error at most $\delta/100L$.

The asserted path s can be taken as the reduced concatenation of the consecutive chain

$$a, b, \bar{a}, \bar{b},$$

where the bar notation stands for orientation reversal. In other words, s is the unique geodesic path in M which is homotopic to the concatenated path relative to the endpoints. We check that the constructed path s satisfies the requirements of Proposition 3.5 as follows.

The first requirement about the length and directions at endpoints of s follows from the length estimate of [LM, Lemma 4.8 (1)].

The second requirement about the parallel transport of $\mathbf{p}$ along s can be checked by considering the effect on the basis vectors. The parallel transport of $\vec{t}_p$ along s ends up $(\delta/10L)$–close to $\vec{t}_p$ by the direction estimates of the first requirement with slightly more careful control of the error. The parallel transport of $\vec{n}_p$ consequentially along a, b, $\bar{a}$, $\bar{b}$ results in four vectors which are $(\delta/10L)$–close to $\vec{n}_p$, $\vec{n}_p(\phi/2)$, $\vec{n}_p(-\phi/2)$, $\vec{n}_p(\phi)$ respectively. This is because the effect of parallel transport of vectors $\vec{u} \perp \vec{t}$ along a or $\bar{a}$ is approximately the reflection about the axis $\mathbb{R}\vec{n}_p$ in the orthogonal complement of $\vec{t}_p$, while along b or $\bar{b}$ the axis is $\mathbb{R}\vec{n}_p(\phi/4)$. Then the parallel transport of $\vec{n}_p$ along the reduced concatenation s can be estimated by the phase estimate of [LM, Lemma 4.8 (2)], (choosing auxiliary framing of the segments a, b, $\bar{a}$, $\bar{b}$ approximately the four vectors above). Combining the estimates for the parallel transport of $\vec{t}_p$ and $\vec{n}_p$ yields the desired estimation for the distance between $\mathbf{q}'$ and $\mathbf{p}(\phi)$.

The third requirement on the homology class of $\hat{s}$ is satisfied because of Lemmas 3.2 and 3.4. In fact, applying those lemmas with δ there taken to be $|10\phi|$, it follows that parallel transport along s takes $\mathbf{p}$ to a frame $\mathbf{q}' \in \mathrm{SO}(M)|_p$ which can be connected to $\mathbf{p}(\phi)$ along a unique shortest path $[\mathbf{q}', \mathbf{p}]$ in $\mathrm{SO}(M)|_p$. Moreover, the parallel-transport path $[\mathbf{p}, \mathbf{q}']$ concatenated with the short path $[\mathbf{q}', \mathbf{p}]$ represents a 1-cycle which is null homologous in $\mathrm{SO}(M)$. Since ϕ is small as chosen, the short path $[\mathbf{q}', \mathbf{p}]$ is homotopic, relative to endpoints, to the concatenation of the shortest path $[\mathbf{q}', \mathbf{p}(\phi)]$ in $\mathrm{SO}(M)|_p$ with the reversal of the frame-spinning path $\hat{\omega} = [\mathbf{p}, \mathbf{p}(\phi)]$.

Therefore, the path s constructed above meets all the requirements of Proposition 3.5, as claimed for the local case.

3.2.2. *The general case.* We proceed to prove the general case of Proposition 3.5. Let $\mathbf{p} = (\vec{t}_p, \vec{n}_p, \vec{t}_p \times \vec{n}_p)$ in $\mathrm{SO}(M)$ be a frame at a point p in M, and let $\phi \in \mathbb{R}$ be an arbitrarily given constant.

Note that replacing ϕ with $\phi + 4k\pi$ for any integer k does not change the homology class of the frame-spinning path $[\hat{\omega}] \in H_1(\mathrm{SO}(M), \mathbf{p} \cup \mathbf{p}(\phi); \mathbb{Z})$, because $\pi_1(\mathrm{SO}(M)|_p) \cong \mathbb{Z}_2$. Hence we may assume that $\phi \in [4\pi, 8\pi)$, and for some universally large integer $D > 0$, say 10^5, we may assume that the value

$$\phi_D = \phi/D > 0$$

is as small as applicable to the local case. In particular, any ball of radius $1000\phi_D$ in $\mathrm{SO}(M)|_p$ is embedded. Given any $\delta > 0$, we may assume that $1000\delta < 4\pi/D$, possibly after replacing it with a smaller value. This ensures that any ball of radius 100δ in $\mathrm{SO}(M)|_p$ is convex and isometrically embedded. Let $L > 1$ be so large that $L_D = L/D$ works for the local case with respect to $\delta_D = \delta/D^2$.

Under the setting above, we construct the asserted path s as follows. By applying the local case, we obtain a geodesic path s_D, so that all the requirements of Proposition 3.5 are satisfied by s_D with respect to the describing parameters (ϕ_D, δ_D, L_D). We take the asserted geodesic path s in M to be the reduced concatenation of the consecutive chain of paths in M:

$$a_1, \cdots, a_D,$$

which consists of D copies a_i of s_D.

The first two requirements of Proposition 3.5 can be shown to be satisfied by s using a similar argument as the local case, which applies the estimates of [LM, Lemma 4.8]. It remains to check that the almost parallel-transport path $\hat{s} \in \pi_1(\mathrm{SO}(M), \mathbf{p}, \mathbf{p}(\phi))$ in the third requirement of Proposition 3.5 is homologous to the frame-spinning path $\hat{\omega}$ in $H_1(\mathrm{SO}(M), \mathbf{p} \cup \mathbf{p}(\phi); \mathbb{Z})$.

To this end, write $\hat{\omega}$ as the concatenation of consecutive subpaths $\hat{\omega}_1, \cdots, \hat{\omega}_D$, where $\hat{\omega}_i$ is the frame-spinning path parametrized by the subinterval $[(i-1)\phi_D, i\phi_D]$ of $[0, \phi]$. The terminal endpoint of $\hat{\omega}_i$ is the frame $\mathbf{p}_i = \mathbf{p}(i\phi_D)$. Parallel transport of $\mathbf{p}$ consequentially along $a_1, \cdots, a_D$ gives rise to a sequence of points $\mathbf{q}_1, \cdots, \mathbf{q}_D$. We also set $\mathbf{q}_0 = \mathbf{p}$. Since each a_i is a copy of s_D, it is easy to estimate that parallel-transport along a_i takes $\mathbf{p}_{i-1}$ to a frame $\mathbf{p}'_i$ in $\mathrm{SO}(M)|_p$, which lies in the $(100\delta_D)$–neighborhood of $\mathbf{p}_i$; and hence, it can be estimated that $\mathbf{q}_i$ lies in the $(100\delta/D)$–neighborhood of $\mathbf{p}_i$. By our assumption, the balls B_i of radius 100δ centered at $\mathbf{p}_i$ are mutually disjointly embedded in $\mathrm{SO}(M)|_p$, and any pair of frames in each B_i can be connected by a unique shortest path in B_i. The fact that a_i is a copy of s_D implies that, as a 1-chain of $\mathrm{SO}(M)$, the parallel-transport path $[\mathbf{q}_{i-1}, \mathbf{q}_i]$ equals $\hat{\omega}_i$ plus some 1-chains in B_{i-1} and B_i, modulo 1-boundaries of $\mathrm{SO}(M)$. It follows that the consequential concatenation of $[\mathbf{q}_{i-1}, \mathbf{q}_i]$, where $i = 0, \cdots, D-1$, is equal to $\hat{\omega}$ plus 1-chains in B_i the modulo 1-boundaries of $\mathrm{SO}(M)$. By Lemma 3.2, the consequential concatenation of the D paths $[\mathbf{q}_{i-1}, \mathbf{q}_i]$ equals the parallel-transport path $[\mathbf{p}, \mathbf{q}']$ of $\mathbf{p}$ along s, up to 1-chains of B_D and 1-boundaries of $\mathrm{SO}(M)$. It follows that the almost parallel-transport path $\hat{s}$ in the third requirement of Proposition 3.5 differs from $\hat{\omega}$ only by 1-boundaries of $\mathrm{SO}(M)$ and 1-cycles of B_i, which are again 1-boundaries of $\mathrm{SO}(M)$. This verifies the third requirement of Proposition 3.5.

Therefore, we have completed the proof of Proposition 3.5 in the general case.

3.3. **Connecting frames in a grading class.** In this section, we prove Theorem 3.1. Suppose that $\mathbf{p} = (\vec{t}_p, \vec{n}_p, \vec{t}_p \times \vec{n}_p)$ and $\mathbf{q} = (\vec{t}_q, \vec{n}_q, \vec{t}_q \times \vec{n}_q)$ in $\mathrm{SO}(M)$ be frames given at points p and q of M, respectively. For any given relative homology class $\Xi \in H_1(\mathrm{SO}(M), \mathbf{p} \cup \mathbf{q}; \mathbb{Z})$, the goal is to construct a geodesic segment s in M from p to q so that the parallel-transportation path of $\mathbf{p}$ along s almost represents Ξ. Furthermore, for any given constant $\delta > 0$, we want s to have length approximately L, and transport $\mathbf{p}$ approximately to $\mathbf{q}$, with error bounded by δ/L, and we want the construction to be possible for any sufficiently large $L > 0$.

The recipe consists of three steps: First, we construct a geodesic segment w from p to q which realizes the projection of Ξ in $H_1(M, p \cup q; \mathbb{Z})$. Then we adjust w by concatenating it nearly smoothly with null homologous closed geodesic paths a, b based at p, q respectively,

so that the reduced concatenation awb have approximately the wanted directions at the endpoints. Finally, we adjust the awb by further concatenating it with a frame-spinning path z based at q, in a nearly smooth fashion, so that the parallel-transport path induced by the reduced concatenation $awbz$ almost represents Ξ. We can appropriately choose the length of the paths a, w, b, z so that their reduced concatenation $awbz$ yields the desired segment s.

To be precise, suppose that a constant $\delta > 0$ is given. Possibly after replacing it with a smaller value (depending the injectivity radius of M), we may assume in addition that any ball of radius 1000δ is embedded in $\mathrm{SO}(M)$. At this point, we simply assume that $L > 1$ is any arbitrary constant which is sufficiently large to enable all the participating constructions and estimates in our recipe. A specific lower bound for L is summarized at the end of our recipe.

The recipe for constructing the asserted s is as follows.

Step 1. We take w to be a geodesic path in M from p to q which realizes the projection of Ξ in $H_1(M, p \cup q; \mathbb{Z})$. Here the projection is induced by the projection map $\mathrm{SO}(M) \to M$. Since all but finitely many such w are long, for any constant $K_0 = K_0(\delta) > 1$, to be specified later, we may assume that the length of w is at least K_0. On the other hand, there exists a constant

$$K_1 = K_1(K_0, M, \Xi) > K_0,$$

so that the length of w is at most K_1, such as the minimal length of paths in the projection of Ξ of length at least $K_0 + 1$. Denote the length of w by

$$K_0 < \ell(w) < K_1.$$

Step 2. We take a to be a geodesic path in M from p to p, and b a geodesic path in M from q to q, so that the following properties hold.

- The length $\ell(a)$ of a is $(\delta/100L)$–close to $L/4 - \ell(w)/2$. The initial direction of a is $(\delta/100L)$–close to $\vec{t}_p$, and the terminal direction of a is $(\delta/100L)$–close to the initial direction of w. Moreover, a is null homologous in M.
- The length $\ell(b)$ of b is $(\delta/100L)$–close to $L/4 - \ell(w)/2$. The initial direction of b is $(\delta/100L)$–close to the terminal direction of w, and the terminal direction of b is $(\delta/100L)$–close to $\vec{t}_q$. Moreover, b is null homologous in M.

The geodesic path a can be taken to be the reduced concatenation of a commutator of two suitably chosen geodesic paths at p, and similarly can we construct b. An explicit construction can be found in [Su2, Lemma 3.4]. For example, to invoke the lemma to create a geodesic path a, one may take the vectors $\vec{v}_1$ and $\vec{v}_2$ there to be $\vec{t}_p$ and the initial direction of w, and the vector $\vec{n}$ there to be any unit vector perpendicular to both $\vec{v}_1$ and $\vec{v}_2$. Let

$$L_1 = L_1(\delta, K_1, M) > 10K_1$$

be an *a priori* lower bound for L to enable the construction for a and b. Note that $\ell(a)$ and $\ell(b)$ are hence at least K_0.

Step 3. We take z to be a geodesic path in M from q to q as follows. Recall that the notation $\mathbf{q}(\phi)$ stands for the rotation of $\mathbf{q}$ about $\vec{t}_q$ by an angle ϕ. By our construction, the parallel transport of $\vec{n}_p$ consequentially along a, w, b is a unit vector $\vec{m}$ at q which is nearly orthogonal to $\vec{t}_q$. Pick a value $\phi \in \mathbb{R}$ such that $\vec{n}_q(-\phi)$ minimizes the distance to $\vec{m}$ among unit vectors orthogonal to $\vec{t}_q$. We apply Proposition 3.5 to construct two candidates of z,

denoted as $z_\uparrow$ and $z_\downarrow$. One of them satisfies the following properties, and the other satisfies the same properties but with ϕ replaced by $\phi + 2\pi$.

- The length of z is $(\delta/100L)$–close to $L/2$. The initial direction and the terminal direction of z are both $(\delta/100L)$–close to $\vec{t}_q$.
- The parallel transport from q back to q along z takes $\mathbf{q}(-\phi)$ to a frame $\mathbf{q}'$ which is $(\delta/100L)$–close to $\mathbf{q}$, and there exists a unique shortest path in $\mathrm{SO}(M)|_p$ between $\mathbf{q}'$ and $\mathbf{q}$.
- Denote by $\hat{z}$ the path which is the concatenation the parallel-transport path from $\mathbf{q}(-\phi)$ to $\mathbf{q}'$ with the shortest path from $\mathbf{q}'$ to $\mathbf{q}$. The relative homology class represented by $\hat{z} \in \pi_1(\mathrm{SO}(M), \mathbf{q}(-\phi), \mathbf{q})$ equals $[\hat{\omega}] \in H_1(\mathrm{SO}(M), \mathbf{q}(-\phi) \cup \mathbf{q}; \mathbb{Z})$.

Note that $\mathbf{q}(-\phi) = \mathbf{q}(-\phi - 2\pi)$ but the corresponding $[\hat{\omega}]$ are different for the two candidates, so the relative homology classes

$$[\hat{z}_\uparrow],\ [\hat{z}_\downarrow] \ \in\ H_1(\mathrm{SO}(M), \mathbf{q}(-\phi) \cup \mathbf{q};\ \mathbb{Z})$$

differ exactly by the canonical element $[\hat{c}] \in H_1(\mathrm{SO}(M); \mathbb{Z})$ of order 2, namely, the nontrivial element of $H_1(\mathrm{SO}(M)|_q; \mathbb{Z}) \cong \mathbb{Z}_2$. Let

$$L_2 = L_2(\delta, K_0, M) > K_0$$

be a lower bound for L so that the application of Proposition 3.5 is valid.

The geodesic path $w, a, b, z_\uparrow, z_\downarrow$ that we have constructed are all longer than the presumed constant $K_0 > 1$. According to [LM, Lemma 4.9] and Lemma 3.2 we can choose $K_0 = K_0(\delta) > 1$ to be sufficiently large, then both of the geodesic paths $s_\uparrow = awbz_\uparrow$ and $s_\downarrow = awbz_\downarrow$, which are obtained by reduced concatenation, satisfy the first two asserted properties of Theorem 3.1. Regarding to the third property, $s_\uparrow$ and $s_\downarrow$ give rise to two paths of frames $\hat{s}_\uparrow$ and $\hat{s}_\downarrow$ from $\mathbf{p}$ to $\mathbf{q}$, and exactly one of them represents $\Xi \in H_1(\mathrm{SO}(M), \mathbf{p} \cup \mathbf{q}; \mathbb{Z})$ while the other represents $\Xi + [\hat{c}]$.

To finish our recipe, we take the correct z and the corresponding reduced concatenation

$$s = awbz$$

so that $\hat{s}$ represents Ξ. Then the geodesic path s in M from p to q meets all the asserted properties of Theorem 3.1.

To summarize, suppose that $\mathbf{p}$, $\mathbf{q}$ in $\mathrm{SO}(M)$ are given. For any constant $\delta > 0$, we choose constants K_0, K_1, L_1, and L_2 in order, and take

$$L_3 = L_3(\delta, M, \Xi)\ = \max(L_1, L_2).$$

Then for any constant $L > L_3$, the asserted geodesic path s in M from p to q of Theorem 3.1 can be construction following the recipe above.

This completes the proof of Theorem 3.1.

3.4. **Proof of Theorem 1.4.** We derive Theorem 1.4 from the relative version Theorem 3.1 as follows.

Given any homology class $\Xi \in H_1(\mathrm{SO}(M); \mathbb{Z})$ and any constant $\epsilon > 0$, we take an arbitrary frame $\mathbf{p} = (\vec{t}_p, \vec{n}_p, \vec{t}_p \times \vec{n}_p) \in \mathrm{SO}(M)|_p$ at a point $p \in M$. Identify $H_1(\mathrm{SO}(M); \mathbb{Z})$ naturally with $H_1(\mathrm{SO}(M), \mathbf{p}; \mathbb{Z})$. Denote by

$$\hat{c} \ \in\ \pi_1(\mathrm{SO}(M))$$

the nontrivial central element, which is unique of order 2. For sufficiently large R with respect to ϵ, M, and $\Xi+[\hat{c}]$, we apply Theorem 3.1 to construct a geodesic path s in M from p to p with the following properties:

- The length of s is $(\epsilon/10R)$–close to R. The initial direction and the terminal directions are both $(\epsilon/10R)$–close to $\vec{t}_p$.
- The parallel transport of $\mathbf{p}$ along s is a frame $\mathbf{p}' \in \mathrm{SO}(M)|_p$ which is $(\epsilon/10R)$–close to $\mathbf{p}$.
- The asserted concatenated path $\hat{s}$ represents $\Xi+[\hat{c}]$ in $H_1(\mathrm{SO}(M), \mathbf{p}; \mathbb{Z})$.

Denote by γ the unique free geodesic loop which is freely homotopic to the closed path s.

It follows from Lemma 3.3 that γ is an $(R, \epsilon/R)$–curve. Moreover, for any frame $\mathbf{q}$ at a point $q \in \gamma$, the cyclic concatenation of the parallel-transport path $[\mathbf{q}, \mathbf{q}']$ of $\mathbf{q}$ around γ with the shortest path $[\mathbf{q}', \mathbf{q}]$ in $\mathrm{SO}(M)|_q$ is freely homotopic to the closed path $\hat{s}$ in $\mathrm{SO}(M)$, but its free homotopy class differs from the canonical lift $\hat{\gamma} \in \pi_1(\mathrm{SO}(M))$ by a factor $\hat{c}$, (see Section 2). Therefore, $[\hat{\gamma}]$ equals Ξ in $H_1(\mathrm{SO}(M); \mathbb{Z})$, as desired.

This completes the proof of Theorem 1.4.

4. Subsurface of odd Euler characteristic

In this section, we construct immersed quasi-Fuchsian subsurfaces of odd Euler characteristic in closed hyperbolic 3-manifolds, proving Theorem 1.1.

Let M be a closed hyperbolic 3-manifold. Without loss of generality, we assume that M is orientable and fix an orientation, otherwise we pass to an orientable double cover. Denote by $\mathrm{SO}(M)$ the special orthonormal frame bundle over M.

Take an arbitrary frame in $\mathrm{SO}(M)$ at a point p in M, denote as

$$\mathbf{p} = (\vec{t}_p, \vec{n}_p, \vec{t}_p \times \vec{n}_p).$$

Denote by $\mathbf{p}^* \in \mathrm{SO}(M)|_p$ the frame

$$\mathbf{p}^* = (\vec{t}_p, -\vec{n}_p, -\vec{t}_p \times \vec{n}_p).$$

For a sufficently small constant $\epsilon > 0$ and some sufficiently large constant $R > 1$, which we specify in the summary at the end of the proof, we apply Theorem 3.1 to construct a geodesic path s in M with both endpoints p satisfying the following requirements:

- The length of s is $(\epsilon/100R)$–close to $R/2$. The initial direction and the terminal direction of s are both $(\epsilon/100R)$–close to $\vec{t}_p$.
- The parallel transport of $\mathbf{p}$ along s back to p is $(\epsilon/100R)$–close to $\mathbf{p}^*$.
- The closed path s in null homologous in $H_1(M; \mathbb{Z})$.

Note that there are two qualified candidates for the relative homology class $\Xi = [\hat{s}]$ in $H_1(\mathrm{SO}(M), \mathbf{p} \cup \mathbf{p}^*; \mathbb{Z})$ as of Theorem 3.1, differing from each other by the homology class of the central nontrivial element $\hat{c} \in \pi_1(\mathrm{SO}(M))$. Either of them works fine. For sufficently small ϵ and large R, the closed geodesic loop of M which is freely homotopic to the cyclic concatenation of two copies of s is a good curve:

$$\gamma \in \mathbf{\Gamma}_{R,\epsilon}(M).$$

The closed geodesic loop of M freely homotopic s itself, denoted as

$$\sqrt{\gamma} \in \pi_1(M),$$

is not good. It is doubly covered by γ, and has complex length approximately $R/2+\pi\cdot\sqrt{-1}$, with error at most $\epsilon/2R$ in absolute value.

For sufficiently small ϵ and large R, we observe the following fact:

Lemma 4.1. *The canonical lift $\hat{\gamma}\in\pi_1(\mathrm{SO}(M))$ is null homologous in $\mathrm{SO}(M)$.*

Proof. Since the element $s\in\pi_1(M,p)$ is homologically trivial, there exists a path of frames $\hat{\alpha}\in\pi_1(\mathrm{SO}(M),\mathbf{p},\mathbf{p}^*)$ which is contained in $\mathrm{SO}(M)|_p$, such that the path $\hat{s}$ from $\mathbf{p}$ to $\mathbf{p}^*$ as in the conclusion of Theorem 3.1 is relatively homologous to $\hat{\alpha}$. Denote by $R_T(\pi)\in\mathrm{SO}(3)$ the matrix

$$R_T(\pi)\ =\ \begin{bmatrix}1&0&0\\0&-1&0\\0&0&-1\end{bmatrix}.$$

It follows that the path

$$\alpha^*=\alpha\cdot R_T(\pi)$$

from $\mathbf{p}^*$ to $\mathbf{p}$ is relatively homologous to $\hat{s}\cdot R_T(\pi)$. By Lemma 3.3 and the definition of the canonical lift, the cyclic concatenation of $\hat{s}$ with $\hat{s}\cdot R_T(\pi)$ differs from $\hat{\gamma}$ by a factor $\hat{c}$, as an element of $\pi_1(\mathrm{SO}(M))$. On the other hand, the concatenation of α with $\hat{\alpha}^*$ represents $\hat{c}$ in $\pi_1(\mathrm{SO}(M))$ by the topology of $\mathrm{SO}(3)$. We conclude that $\hat{\gamma}$ is homologically $2[\hat{c}]=0$ in $H_1(\mathrm{SO}(M);\mathbb{Z})$. □

Therefore, it follows from Theorem 2.4 that there exists a π_1-injectively immersed, (R,ϵ)–panted, connected, quasi-Fuchsian subsurface

$$F\looparrowright M$$

which is oriented and bounded by γ. In fact, the complex Fenchel–Nielsen coordinates associated to any glued cuff C of F can be required to satisfy $(\mathbf{hl}(C),s(C))\approx(R/2,1)$ with error at most $(\epsilon/R,\epsilon/R^2)$ componentwise, the same as in the original construction of Kahn and Markovic (Theorem 2.2 and Declaration 2.5).

We would like to identify every pair of antipodal points of γ, or topologically to glue to the boundary of F a Möbius band with the core $\sqrt{\gamma}$, so as to produce an immersed non-orientable subsurface Σ as our output. However, to guarantee that Σ is π_1–injective, we want to make sure that geometrically F contains no properly embedded essential arcs which are relatively short compared to R. This is the last technical point that we need to address before completing the construction.

To this end, denote by $\mathcal{G}(F)$ the dual graph of the inherited pants decomposition of F. Namely, the vertices of $\mathcal{G}(F)$ are the (R,ϵ)–pants of F, and the edges of $\mathcal{G}(F)$ are the glued cuffs, which are the (R,ϵ)–curves in the interior of F. The graph $\mathcal{G}(F)$ is trivalent except at the distinguished valence-2 vertex P_0 which contains ∂F.

Lemma 4.2. *The panted subsurface F can be constructed to satisfy the extra condition that every non-contractible closed combinatorial path of $\mathcal{G}(F)$ based at P_0 has combinatorial length at least $Re^{R/4}$.*

Proof. Suppose that F_0 is an oriented connected (R,ϵ)–panted subsurface with connected boundary γ, as guaranteed by Theorem 2.4. We may require F_0 to have complex Fenchel–Nielsen coordinates $(\epsilon/R,\epsilon/R^2)$–close to $(R/2,1)$ for every glued cuff. See [Su2, Corollary 2.7] for an outline of the construction based on [LM].

Denote by P_0 the distinguished pair of pants of F_0 with one cuff ∂F. Denote by C_1, C_2 the other two cuffs of P_0, and P_1, P_2 the other two pairs of pants adjacent to P_0 along C_1, C_2 accordingly. Take an oriented connected closed (R, ϵ)–panted surface E with complex Fenchel–Nielsen coordinates $(\epsilon/R, \epsilon/R^2)$–close to $(R/2, 1)$, (Theorem 2.2). We may require that (the pants types of) P_1 and P_2 also appear in the inherited pants decomposition of E. This follows from [LM, Theorems 2.9 and 2.10]. Possibly after passing to a finite cover of E induced by a regular finite cover of the dual graph $\mathcal{G}(E)$, we may assume that every embedded cycle of $\mathcal{G}(E)$ has combinatorial length at least $Re^{R/4}$ and no edge of $\mathcal{G}(E)$ is separating.

We modify F_0 to obtain a new subsurface F with the asserted property. Take a copy E' of E such that some pair of pants $P_1' \subset E'$ has the same pants type of P_1, with a cuff C_1' corresponding to C_1. Denote by $P' \subset E'$ the pants adjacent to P_1' along C_1'. We make a cross change between F_0 and E' along the parallel glued cuffs C_1 and C_1'. Namely, cut F_0 along C_1, and E' along C_1'; identify the new unglued cuff of P_1' with the new unglued cuff of P_0, and similarly glue the pants P_1 and P' along their new unglued cuffs. In the same way, we make a cross change between F_0 and another copy E'' of E along cuffs corresponding to C_2. In effect, we obtain a new (R, ϵ)–panted subsurface F which is oriented and bounded by γ. The complex Fenchel–Nielsen coordinates remain $(\epsilon/R, \epsilon/R^2)$–close to $(R, 1)$. The pants decomposition graph $\mathcal{G}(F)$ is obtained from $\mathcal{G}(F_0)$ and two copies of $\mathcal{G}(E)$ by two cross changes of edges corresponding to C_1 and C_2. It is straightforward to see that $\mathcal{G}(F)$ has no non-contractible closed path based at P_0 of combinatorial length smaller than $Re^{R/4}$. Therefore, the panted subsurface F is as desired. □

Take F to be a connected oriented (R, ϵ)–panted immersed subsurface of M bounded by the (R, ϵ)–curve γ, such that the complex Fenchel–Nielsen coordinates are approximately $(R/2, 1)$ for every glued cuff, with error at most $(\epsilon/R, \epsilon/R^2)$ componentwise. Suppose in addition that the pants decomposition graph $\mathcal{G}(F)$ satisfies the conclusion of Lemma 4.2. Since γ doubly cover the closed geodesic $\sqrt{\gamma}$ of M, the pre-image of any point of $\sqrt{\gamma}$ is a pair of antipodal points in γ. We identify every pair of antipodal points of ∂F accordingly.

The result is a connected closed immersed subsurface

$$\Sigma \looparrowright M,$$

which is no longer orientable. However, it is π_1–injectively immersed and geometrically finite, hence quasi–Fuchsian, by a geometric criterion due to H. Sun [Su1, Theorem 2.6]. Furthermore, the Euler characteristic of S can be computed by:

$$\chi(\Sigma) \,=\, \chi(F) \,= 1 - 2 \cdot \mathrm{genus}(F),$$

which is an odd number.

In summary, we can choose some sufficiently small $\epsilon > 0$ according to M and some sufficiently large $R > 0$ according to M and ϵ, and construct a closed quasi–Fuchsian subsurface $\Sigma \looparrowright M$ of odd Euler characteristic as asserted. Fixing a choice of ϵ, it suffices to require that R should be so large that all the constructions and estimates that we have done work.

This completes the proof of Theorem 1.1.

5. Irregular exhausting tower with exponential torsion growth

In this section, we present the core construction of Theorem 1.3, which can be stated in slightly more details as follows.

Proposition 5.1. *Let M be an orientable closed hyperbolic 3-manifold. Suppose that there are closed surfaces S and Σ, and there are maps $\iota_S\colon S \to M$, and $\iota_\Sigma\colon \Sigma \to M$, and $\rho\colon M \to S$ with the following properties:*

- *The surface S is orientable, and Σ of odd Euler characteristic.*
- *The maps ι_S and ι_Σ are embeddings with mutually disjoint images.*
- *The composition $\rho \circ \iota_S$ is homotopic to the identity, and the composition $\rho \circ \iota_\Sigma$ is homotopic to a constant point map.*

Then for any constant $0 < \epsilon < 1$ and any base point $*$ of M, there exists a tower of finite covers of M with distinguished lifted base points:

$$\cdots \longrightarrow (\tilde{M}_n, \tilde{o}_n) \longrightarrow \cdots \longrightarrow (\tilde{M}_2, \tilde{o}_2) \longrightarrow (\tilde{M}_1, \tilde{o}_1) \longrightarrow (M, *).$$

Moreover, the following requirements are satisfied:

- *The injectivity radii of $\tilde{M}_n$ at the base points are unbounded.*
- *The number of distinct lifts of ι_Σ into each $\tilde{M}_n$ is at least $(1-\epsilon)\cdot[\tilde{M}_n : M]$. Hence,*

$$\lim_{n\to\infty} \frac{\log|\mathrm{Tor}_1(H_1(\tilde{M}_n;\mathbb{Z}),\mathbb{Z}_2)|}{[\tilde{M}_n : M]} \geq (1-\epsilon)\log 2.$$

To illustrate the strategy, let us briefly explain the first move. That is, given a nontrivial element $g_1 \in \pi(M, *)$, how to construct a finite cover $(\tilde{M}_1, \tilde{o}_1)$ into which g_1 does not lift but many Σ lifts. We can first take a cyclic cover (M_1', o_1') dual to S of some large degree d. Writing V for the complement of S in M, we can decompose M_1' into d lifted copies V_i of V, where $i \in \mathbb{Z}_d$. If g_1 intersects S algebraically nontrivially, we can simply take d to be very large and $\tilde{M}_1$ to be M_1'. Otherwise, g_1 lifts into M_1' based at o_1', and it is contained in the union W of some consecutive pieces V_i nearby. As the number of pieces in W depends only on g_1, W occupies a very small portion of M_1' if we choose d to be large. To construct $\tilde{M}_1$, the idea is to assemble a finite cover of W and some finite covers of other pieces V_j. We want to require that g_1 does not lift to the finite cover of W, and meanwhile, that every preimage component of Σ in the finite cover of any other piece V_j is a lift of Σ. The latter can be ensured if the finite cover of V_j comes from a finite cover of S, via the inclusion of V into M and the retraction of M onto S. The former can be ensured by the residual finiteness of W. However, in order to glue these individual covering pieces together, it is crucial to know that their restriction to the boundary are isomorphic covers of S, better characteristic. This is guaranteed by the so-called omnipotence lemma (Lemma 5.3), a consequence of Wise's Malnormal Special Quotient Theorem. Assume that all these are done, then we will obtain a desired $\tilde{M}_1$ in which g_1 disappears but most lifts of Σ survive.

The construction can be iterated to destroy more elements from $\pi_1(M, *)$, one at each time. If the degrees of cyclic covers are chosen to grow very fast, the total portion of non-lifts (or 'damaged lifts') of Σ will remain small. To argue by induction in a formal way, we introduce some notion called *generalized digital expansion* to encode a cyclic covering tower $\{M_n'\}$ underlying the asserted $\{\tilde{M}_n\}$. A precise induction hypothesis can be found as the statement of Lemma 5.4.

The rest of this section is devoted to the proof of Proposition 5.1.

5.1. **Cyclic towers encoded by generalized digital expansions.** In this subsection, we introduce a tower of finite cyclic covers which is to be considered as an intermediate step toward the construction of the exhausting tower asserted by Proposition 5.1.

For this subsection, we suppose that $(M, *)$ is an orientable closed 3-manifold, and S is an embedded non-separating oriented connected closed subsurface of M which misses $*$. Denote by

$$V \;=\; M \setminus \mathrm{Nhd}^{\circ}(S)$$

the compact submanifold of M obtained by cutting along S. Let d be an odd positive integer which is at least 3.

5.1.1. *Generalized digital expansion and blocks.* For any positive integer n, denote by $[d^n]$ the set of the d^n consecutive integers centered at 0, namely,

$$[d^n] \;=\; \left\{0, \pm 1, \pm 2, \cdots, \pm \frac{d^n - 1}{2}\right\}.$$

Adopt the notation

$$s_n = 1 + 2 + \cdots + n = \frac{n(n+1)}{2}.$$

There is a canonical bijective correspondence between sets:

$$\begin{array}{ccc} [d^{s_n}] & \longleftrightarrow & [d] \times [d^2] \times \cdots \times [d^n] \\ a & \leftrightarrow & (a_0,\, a_1,\, \cdots,\, a_{n-1}) \end{array}$$

which is determined by the relation

$$a \;=\; \sum_{i=0}^{n-1} a_j d^{s_j}.$$

We view the correspondence as a generalized digital expansion for the integers a of $[d^{s_n}]$: At the j-th place with the assigned weight d^{s_j}, the digit a_j is taken from the digit set $[d^{j+1}]$, which is particular for that place.

We say that two integers $a, a' \in [d^{s_n}]$ are in *the same block*, denoted as

$$a \sim a',$$

if the generalized digital expansions of a and a' agree from the $(n-1)$–th place all the way down to the highest place with a digit 0. In other words, if $a_j \neq 0$ for all $j = 0, \cdots, n-1$, then $a \sim a'$ means $a_j = a'_j$ for all $j = 0, \cdots, n-1$; if $a_k = 0$ for some $k \in \{0, \cdots, n-1\}$ and $a_j \neq 0$ for all $j = k+1, \cdots, n-1$, then $a \sim a'$ means $a_j = a'_j$ for all $j = k, \cdots, n-1$. Since being in the same block is an equivalence relation on $[d^{s_n}]$, we call the equivalence classes the *blocks* of $[d^{s_n}]$. Denote the set of blocks of $[d^{s_n}]$ as

$$\mathscr{B}_n = \mathscr{B}_n(d) = [d^{s_n}]/\sim .$$

For any block $\beta \in \mathscr{B}_n$, the *level* of β is said to be $k \in \{0, \cdots, n-1\}$, if for some (hence any) integer $a \in \beta$, a_k is the highest 0 in the expansion of a; the level of β is formally defined to be ∞, if the block β consists of a single integer which has no 0 digits.

For example, the subset $[d^{s_{n-1}}]$ of $[d^{s_n}]$ is the only block of level $(n-1)$. Its size is $1/d^n$ of $[d^{s_n}]$. Any of the remaining blocks of $[d^{s_n}]$, say of a level j other than ∞, looks like a shifted subset $[d^{s_j}]$ centered at some integer with an expansion $(0, \cdots, 0, a_{j+1}, \cdots, a_{n-1})$, where $a_{j+1}, \cdots, a_{n-1}$ are nonzero. For large n, the blocks of level ∞ in $[d^{s_n}]$ reminds us of the picture of Cantor's dust set, but the dust is really heavy:

Lemma 5.2. *Given any constant $0 < \epsilon < 1$, the following statement holds true for all sufficiently large odd positive integers d: For all $n \in \mathbb{N}$, the number of blocks of level ∞ in $[d^{s_n}]$ is greater than $(1-\epsilon)\cdot d^{s_n}$.*

Proof. The number C_n of the blocks of level ∞ in $[d^{s_n}]$ is clearly $(d-1)\times\cdots\times(d^n-1)$. For an odd positive integer d at least 3, the portion of such blocks $(C_n \,/\, d^{s_n})$ is strictly decreasing as n grows. The limit can be expressed as $\phi(1/d)$ using the Euler function:

$$\begin{aligned}\phi(q) &= \prod_{j=1}^{\infty}(1-q^n) = \sum_{k=-\infty}^{\infty}(-1)^k q^{(3k^2-k)/2}\\ &= 1-q^2-q^3+q^5+q^7-\cdots\end{aligned}$$

When d is sufficiently large, we have $\phi(1/d) > 1-\epsilon$, so C_n is at least $(1-\epsilon)\cdot d^{s_n}$. □

5.1.2. *Encoding a tower of finite cyclic covers.* Continue to adopt the notations of the generalized digital expansion with respect to d. Denote by

$$\cdots \longrightarrow M'_n \longrightarrow \cdots \longrightarrow M'_2 \longrightarrow M'_1 \longrightarrow M$$

the tower of finite cyclic covers dual to S with the covering degree

$$[M'_n : M] = d^{s_n},$$

so M'_n covers M'_{n-1} cyclically of degree d^n.

Since V is the compact submanifold obtained by cutting M along S, the boundary ∂V has two components $\partial_\pm V$ parallel to S inside M such that the (outward) induced orientation of $\partial_\pm V$ coincides with $\mp S$. The cyclic cover M'_n can be constructed by gluing d^{s_n} copies $V_n(a)$ of V in such a way that $\partial_+ V_n(a)$ is identified with $\partial_- V_n(a+1)$ for all $a \in [d^{s_n}]$. By convention, $V_n(\frac{d^{s_n}-1}{2}+1)$ stands for $V_n(-\frac{d^{s_n}-1}{2})$. Moreover, each block $\beta \in \mathscr{B}_n$ corresponds to a *block of pieces* $W_n(\beta)$ of M'_n obtained by gluing the *unit pieces* $V_n(a)$ for all $a \in \beta$.

Therefore, we have the following decompositions of M'_n by various lifts of S:

$$M'_n = \bigcup_{a\in[d^{s_n}]} V_n(a) = \bigcup_{\beta\in\mathscr{B}_n} W_n(\beta).$$

Equip each M'_n with the lifted base point o'_n which lies in the unit piece $V_n(0)$. Therefore, o'_n is contained in the unique block of pieces,

$$o'_n \in W_n(\beta_0) \subset M'_n,$$

where $\beta_0 \in \mathscr{B}_n$ is the unique block of level $(n-1)$; for the blocks $\beta \in \mathscr{B}_n$ of level ∞, the corresponding $W_n(\beta)$ are all isomorphic to V.

5.1.3. *Boundary-characteristic finite covers for blocks of pieces.* In literature, [DLW, PW] for example, boundary-characteristic finite covers have been considered for JSJ pieces of irreducible closed 3-manifolds to construct interesting finite covers. We consider a very similar situation where V and S play the roles of a JSJ piece and a JSJ torus accordingly. Continue to consider the tower of finite cyclic covers of M encoded by the generalized digital expansion with respect to d.

For any block of pieces $W_n(\beta)$, we say that a (possibly disconnected) finite cover $\tilde{\mathcal{W}}$ of $W_n(\beta)$ is $\tilde{S}$*–boundary-characteristic*, if $\tilde{S}$ is a characteristic finite cover of S and if every boundary component of $\tilde{\mathcal{W}}$ is isomorphic to $\tilde{S}$, as a cover of S given by the composition $\tilde{\mathcal{W}} \to W_n(\beta) \to M$. Recall that a characteristic cover $\tilde{X}$ of a path-connected space X is a covering space of X which corresponds to a characteristic subgroup of $\pi_1(X)$, namely, a subgroup

invariant under action of the automorphism group $\mathrm{Aut}(\pi_1(X))$. If $\tilde{X}$ is a characteristic cover of some characteristic cover of X, it is also characteristic over X.

The following omnipotence lemma is a consequence of the Malnormal Special Quotient Theorem due to D. T. Wise [Wi, Theorem 12.3], (see [AGM] for an alternate proof). This lemma allows us to produce boundary-characteristic finite covers as deep as we wish. The assumption can certainly be replaced by a weaker one that M is hyperbolic and V has incompressible acylindrical boundary.

Lemma 5.3. *If M is hyperbolic and S is a retract of M, then for every finite cover $\mathcal{W}''$ of a block of pieces $W_n(\beta)$, there exists a characteristic finite cover $\tilde{S}^*$ of S such that the following statement holds true:*

For every characteristic finite cover $\tilde{S}$ of $\tilde{S}^*$, there exists a regular finite cover $\tilde{\mathcal{W}}$ of $W_n(\beta)$ which is $\tilde{S}$–boundary-characteristic. Moreover, the covering projection of $\tilde{\mathcal{W}}$ to $W_n(\beta)$ factors through $\mathcal{W}''$.

Proof. Observe that it suffices to prove for any connected $\mathcal{W}''$, otherwise taking $\tilde{S}^*$ to be a common characteristic finite cover of those constructed componentwise. If $\beta \in \mathscr{B}_n$ is a block of level j other than ∞, the block of pieces $W_n(\beta)$ is homeomorphic to M'_j removing a lift of S. Therefore, it suffices to prove for $W_n(\beta) \cong V$, otherwise arguing using M'_j instead of M. We may also assume that any given $\mathcal{W}''$ is characteristic over V, otherwise replacing it with a further one such.

Because M is atoroidal and S is a retract of M, the inclusion of S into M induces an embedding of $H = \pi_1(S)$ into $\pi_1(M)$ is as (a representative of the conjugacy class of) a malnormal subgroup. Accordingly, the peripheral subgroups $H_\pm = \pi_1(\partial_\pm V)$ of the word hyperbolic group $G = \pi_1(V)$ form a malnormal pair of quasi-convex subgroups. By [Wi, Theorem 16.6], G is virtually special. Moreover, it is implied by Wise's Malnormal Special Quotient Theorem [Wi, Theorem 12.3] that there exist finite-index subgroups $\tilde{H}^*_\pm$ of $H_\pm$ with the following property: For any further finite-index subgroup $\tilde{H}_\pm$ of $\tilde{H}^*_\pm$, the quotient $G \,/\, \langle\langle \tilde{H}_+, \tilde{H}_- \rangle\rangle$, of G by the normal closure of $\tilde{H}_\pm$, is word hyperbolic and virtually special. In particular, it is residually finite. Without loss of generality, we may assume that $\tilde{H}^*_\pm$ are chosen to be characteristic in $H_\pm$, and isomorphic to the same characteristic finite-index subgroup $\tilde{H}^*$ of H. Since the given $G'' = \pi_1(\mathcal{W}'')$ is characteristic as we have assumed in addition, we may also assume that $\tilde{H}^*$ is chosen so deep that any conjugate of $H_\pm$ intersects G'' in exactly the corresponding conjugate of $\tilde{H}^*_\pm$. Finally, we take the asserted characteristic finite cover $\tilde{S}^*$ of S to be the one corresponding to $\tilde{H}^*$.

To verify the stated property, for any characteristic cover $\tilde{S}$ of $\tilde{S}^*$, denote by $\tilde{H}_\pm \cong \tilde{H} = \pi_1(\tilde{S})$ the corresponding subgroup of $\tilde{H}^*_\pm \cong \tilde{H}^* = \pi_1(\tilde{S}^*)$. Using the residual finiteness of $G \,/\, \langle\langle \tilde{H}_+, \tilde{H}_- \rangle\rangle$, we can find a finite-index normal subgroup $\tilde{G}$ of G which intersects $H_\pm$ in exactly $\tilde{H}_\pm$. Take $\tilde{G}''$ to be $\tilde{G} \cap G''$. Then any conjugate of $\tilde{H}_\pm$ intersects $\tilde{G}''$ in exactly the corresponding conjugate of $\tilde{H}_\pm$ as well. Therefore, the regular finite cover $\tilde{\mathcal{W}}$ of V corresponding to $\tilde{G}$ has every boundary component isomorphic to $\tilde{S}$, and the covering projection $\tilde{\mathcal{W}} \to V$ factors through the given intermediate cover $\mathcal{W}''$. □

5.2. **Construction of the asserted tower.** We construct the asserted tower of Proposition 5.1 adopting the notations and assumptions there.

Given a constant $0 < \epsilon < 1$, choose a sufficiently large odd positive integer d as provided by Lemma 5.2. The notations such as s_n, $[d^n]$, $\mathscr{B}_n$ from the generalized digital expansion

with respect to d remains effective for the rest of this section. Note that the existence of the retract ρ forces the oriented connected subsurface S to be non-separating. Denote by

$$\cdots \longrightarrow (M'_n, o'_n) \longrightarrow \cdots \longrightarrow (M'_2, o'_2) \longrightarrow (M'_1, o'_1) \longrightarrow (M, *)$$

the tower of finite cyclic covers dual to S, of degree d^{s_n} over M, and with lifted base points, as before.

Take a sequence which includes all the elements of $\pi_1(M, *)$, denoted by

$$\{\, g_n \in \pi_1(M, *) \,\}_{n\in\mathbb{N}}.$$

By inserting trivial elements between terms, we may assume that the sequence satisfies the following additional properties:

- The first element g_1 is trivial.
- For every $n \in \mathbb{N}$, the algebraic intersection number $\langle [g_n], [S] \rangle$ is bounded strictly by d^{s_n} in absolute value. Here $[g_n] \in H_1(M; \mathbb{Z})$ and $[S] \in H_2(M; \mathbb{Z})$ are the homology classes accordingly.
- Furthermore, if $\langle [g_n], [S] \rangle$ equals 0, then the based lift of g_n in M'_n is contained in the base block of pieces $W_n(\beta_0)$, up to based homotopy.

Under the above setting, the asserted tower can be constructed by the following lemma. In the context of coverings spaces, the term *elevation* is customarily used to mean a preimage component of a sub-manifold in the referred cover, so as to distinguish from the more common term *lift*, which is equivalently a homeomorphic elevation.

Lemma 5.4. *Under the assumptions of Proposition 5.1 and with the notations above, there are finite connected characteristic covers $\{\tilde{S}_n \to S\}_{n\in\mathbb{N}}$ and base-pointed connected finite (irregular) covers $\{\, (\tilde{M}_n, \tilde{o}_n) \to (M'_n, o'_n) \,\}_{n\in\mathbb{N}}$ such that the following properties are satisfied:*

- *Each $\tilde{M}_{n+1}$ is a finite cover of $\tilde{M}_n$, and the covering maps fit into the commutative diagram:*

$$\begin{array}{ccc} \tilde{M}_{n+1} & \longrightarrow & \tilde{M}_n \\ \big\downarrow & & \big\downarrow \\ M'_{n+1} & \longrightarrow & M'_n \end{array}$$

- *The element g_n does not lift into $\tilde{M}_n$ based at $\tilde{o}_n$ for any $n \in \mathbb{N}$, unless it is trivial.*
- *For each block $\beta \in \mathscr{B}_n$ of level ∞ and for every elevation $\tilde{W}$ in $\tilde{M}_n$ of $W_n(\beta) \cong V$, the induced covering map $\tilde{W} \to V$ is isomorphic to the pull-back of the covering map $\tilde{S}_n \to S$ via the composition:*

$$V \xrightarrow{\text{incl.}} M \xrightarrow{\rho} S.$$

Proof. We construct $\tilde{S}_n$ and $(\tilde{M}_n, \tilde{o}_n)$ by induction on $n \in \mathbb{N}$. For n equal to 1, we can simply take $\tilde{S}_1$ to be S and $(\tilde{M}_1, \tilde{o}_1)$ to be (M'_1, o'_1). Suppose that all its previous stages have been completed, then we proceed with the n-th stage, where n is greater than 1.

We introduce a few notations. Denote

$$(\mathcal{M}^\times_n,\, o^\times_n) \;=\; \left(M'_n \times_{M'_{n-1}} \tilde{M}_{n-1},\, (o'_n, \tilde{o}_{n-1})\right),$$

the fiber product of base-pointed covering spaces. The space $\mathcal{M}^\times_n$ is a possibly disconnected base-pointed finite cover of M'_{n-1} that factors through both M'_n and $\tilde{M}_{n-1}$. It can be concretely described as follows: As d is an odd number, there is a unique lift S'_{n-1} of S in M'_{n-1}

furthermost from o'_{n-1}; therefore, $\mathcal{M}_n^\times$ can be obtained as a cyclic cover of $\tilde{M}_{n-1}$ of degree d^n, which is dual to the preimage of S'_{n-1}. For any block $\beta \in \mathscr{B}_n$, denote by $\mathcal{W}_n^\times(\beta)$ the preimage of the block of pieces $W_n(\beta) \subset M'_n$. Observe that $\mathcal{W}_n^\times(\beta)$ are all $\tilde{S}_{n-1}$–boundary-characteristic over $W_n(\beta)$. In fact, by the induction hypothesis, for any block $\gamma \in \mathscr{B}_{n-1}$ of level ∞, every elevation $\tilde{W}$ of $W_{n-1}(\gamma)$ in $\tilde{M}_{n-1}$ is all $\tilde{S}_{n-1}$–boundary-characteristic. Then the observation follows as every component of $\mathcal{W}_n^\times(\beta)$ is next to a lift of such $\tilde{W}$. Denote the disjoint union of all $\mathcal{W}_n^\times(\beta)$ by

$$\mathcal{W}_n^\times \;=\; \bigsqcup_{\beta\in\mathscr{B}_n} \mathcal{W}_n^\times(\beta).$$

The gluing is formally given by an orientation-reversing free involution:

$$\partial\mathcal{W}_n^\times \;\xrightarrow{\;\nu^\times\;}\; \partial\mathcal{W}_n^\times$$

which commutes with the underlying gluing identification between the components of all $\partial W_n(\beta)$ via the covering projection. Identifying each orbit of $\nu_n^\times$ to a point yields

$$\mathcal{M}_n^\times \;=\; \mathcal{W}_n^\times \,/\, \nu_n^\times.$$

There are two simple cases where the construction of $\tilde{S}_n$ and $(\tilde{M}_n, \tilde{o}_n)$ is straightforward. If the element $g_n \in \pi_1(M, *)$ is trivial, $\tilde{S}_n$ can be taken simply as $\tilde{S}_{n-1}$, and $(\tilde{M}_n, \tilde{o}_n)$ can be taken as the base component of $(\mathcal{M}_n^\times, o_n^\times)$. Similarly, if the algebraic intersection number $\langle [g_n], [S]\rangle$ in M is nontrivial, hence less than d^{s_n} in absolute value. By the assumptions on $\{g_n\}_{n\in\mathbb{N}}$, the element g_n does not lift to M'_n or $\mathcal{M}_n^\times$. So again, we can take $\tilde{S}_n$ and $(\tilde{M}_n, \tilde{o}_n)$ the same as the trivial case.

It remains to prove the essential case when g_n is nontrivial with $\langle [g_n], [S]\rangle = 0$ in M. In this case, g_n lifts to be a nontrivial element of $\pi_1(M'_n, o'_n)$. Moreover, abusing the notation, we have

$$g_n \in \pi_1(W_n(\beta_0), o'_n)$$

by the assumption of $\{g_n\}_{n\in\mathbb{N}}$. Denote by $\mathcal{W}_n^\times(\beta_0)$ the preimage of $W_n(\beta_0)$ in $\mathcal{M}_n^\times$. By the residual finiteness of $\pi_1(W_n(\beta_0), o'_n)$, there exists a further regular finite cover $\mathcal{W}''_n(\beta_0)$ of $W_n(\beta_0)$ which factors through $\mathcal{W}_n^\times(\beta_0)$, such that g_n does not lift to $\mathcal{W}''_n(\beta_0)$. By choosing any base point o''_n lifting $o_n^\times$, we have

$$g_n \notin \pi_1(\mathcal{W}''_n(\beta_0), o''_n).$$

Note that the furthermost lift S'_{n-1} of S in M'_{n-1} above has the property that its complement lifts into M'_n as the interior of $W_n(\beta_0)$. Moreover, S'_{n-1} is a retract of M'_{n-1}, since S is a retract of M as assumed by Proposition 5.1. By applying Lemma 5.3 with respect to M'_{n-1}, S'_{n-1} and $\mathcal{W}''_n(\beta_0)$, we obtain a finite characteristic cover $\tilde{S}^*_n(\beta_0)$ of $S'_{n-1} \cong S$. It has the property that for any further finite characteristic cover of S, boundary-characteristic finite covers of $\mathcal{W}''_{n-1}(\beta_0)$ with the prescribed boundary pattern exist, and can be constructed to factor through $\mathcal{W}''_n(\beta_0)$. Similarly, for each block $\beta \in \mathscr{B}_n$ of level j other than ∞ or $(n-1)$, apply Lemma 5.3 with respect to M'_j, S'_j and $\mathcal{W}_n^\times(\beta)$ to obtain a finite characteristic cover $\tilde{S}^*_n(\beta)$ of S; for each $\beta \in \mathscr{B}$ of level ∞, we can simply take $\tilde{S}^*_n(\beta)$ to be $\tilde{S}_{n-1}$. We take a finite characteristic cover

$$\tilde{S}^*_n \longrightarrow S$$

which factors through $\tilde{S}^*_n(\beta)$ for all $\beta \in \mathscr{B}_n$. Note that $\tilde{S}^*_n \to S$ factors through $\tilde{S}_{n-1}$.

To construct the claimed $\tilde{S}_n$ and $(\tilde{M}_n, \tilde{o}_n)$, we first take the claimed finite characteristic cover

$$\tilde{S}_n \longrightarrow S$$

to be a finite characteristic cover of $\tilde{S}_{n-1}$ which factors through $\tilde{S}_n^*$. The claimed base-pointed finite cover $(\tilde{M}_n, \tilde{o}_n)$ can be constructed by merging $\tilde{S}_n$–boundary-characteristic covers of $\mathcal{W}_n^\times(\beta)$ as follows.

For each block $\beta \in \mathscr{B}_n$, we first construct an $\tilde{S}_n$–boundary-characteristic cover $\tilde{\mathcal{W}}_n(\beta)$ of $\mathcal{W}_n^\times(\beta)$, in the following way. For the block $\beta_0 \in \mathscr{B}_n$ of level $(n-1)$, take an $\tilde{S}_n$–boundary-characteristic finite cover of $\mathcal{W}_n^\times(\beta_0)$ which factors through $\mathcal{W}_n''(\beta_0)$, denoted as $\tilde{\mathcal{W}}_n(\beta_0)$. Similarly, for each block $\beta \in \mathscr{B}_n$ of level j other than ∞ or $(n-1)$, take an $\tilde{S}_n$–boundary-characteristic finite cover of $\mathcal{W}_n^\times(\beta)$. For each block $\beta \in \mathscr{B}_n$ of level ∞, we need a more specific construction of $\tilde{\mathcal{W}}_n(\beta)$ to meet the third property of Lemma 5.4. Note that $\mathcal{W}_n^\times(\beta)$ is isomorphic to $\tilde{\mathcal{W}}_{n-1}(\gamma) \times [d^n]$ for some level–∞ block $\gamma \in \mathscr{B}_{n-1}$. Therefore, by the induction hypothesis, any component $W^\times$ of $\mathcal{W}_n^\times(\beta)$, as a covering space of $W_n(\beta) \cong V$, is isomorphic to the pull-back via $\tilde{S}_{n-1} \to S$ of

$$V \xrightarrow{\rho_V} S,$$

the composition of the retraction ρ of M to S with the inclusion of V into M. In other words, it is isomorphic to a fiber product:

$$W^\times \cong V \times_{\rho_V} \tilde{S}_{n-1}.$$

Denote the $\tilde{S}_n$–boundary-characteristic cover of $W_n(\beta) \cong V$:

$$\tilde{W} \cong V \times_{\rho_V} \tilde{S}_n,$$

then $\tilde{W}$ is isomorphic to the pull-back of $\tilde{S}_n$ via ρ_V and the covering projection factors through $W^\times$. As $W^\times$ runs over the components of $\mathcal{W}^\times(\beta)$, take $\tilde{\mathcal{W}}_n(\beta)$ for the level–∞ block β to be the disjoint union of all the $\tilde{W}$ accordingly. Let us formally put together all the building parts $\tilde{\mathcal{W}}_n(\beta)$ that we have constructed above:

$$\tilde{\mathcal{W}}_n = \bigsqcup_{\beta \in \mathscr{B}_n} \tilde{\mathcal{W}}_n(\beta),$$

At this point, we may not yet be able to obtain a closed manifold by assembling the components of $\tilde{\mathcal{W}}_n$, because the boundary components of $\tilde{\mathcal{W}}_n$ that we wanted to glue up together may not be balanced in amount between opposite orientations. We need to suitably duplicate the components of $\tilde{\mathcal{W}}_n$ to meet the balance condition. A simple solution is to introduce a quantity

$$K_n = K_n(\tilde{\mathcal{W}}_n \to \mathcal{W}^\times)$$

which is defined to be the least common multiple of all the local covering degrees $[\tilde{\mathcal{W}}_n : \mathcal{W}_n^\times]_{W^\times}$ where $W^\times$ runs over all the components of $\mathcal{W}_n^\times$. Here the (unsigned) local covering degree $[\tilde{\mathcal{W}}_n : \mathcal{W}_n^\times]_{W^\times}$ is defined to be the number of lifts in $\tilde{\mathcal{W}}_n$ for any point of $W^\times$. Replace the preimage of each component $W^\times$ by the disjoint union of $K_n / [\tilde{\mathcal{W}}_n : \mathcal{W}_n^\times]_{W^\times}$ copies of itself. It is easy to see that the new cover $\tilde{\mathcal{W}}_n$ has constant local degree $[\tilde{\mathcal{W}}_n : \mathcal{W}_n^\times]_{W^\times} = K_n$ for all $W^\times$. So the balance condition is satisfied as $\tilde{\mathcal{W}}_n$ is already boundary-characteristic over $\mathcal{W}_n^\times$.

We can construct the claimed $\tilde{M}_n$ by gluing up the components of $\tilde{\mathcal{W}}_n$ along boundary. There is a fairly routine procedure to do so. We provide some details below for the reader's reference. (See [DLW, PW] for similar constructions with respect to JSJ decompositions.)

The balance condition allows us to construct a free involution $\tilde{\nu}_\sharp$ which pairs up oppositely oriented boundary components and commutes with the pairing free involution $\nu_\sharp^\times$ via the covering projection, namely, the following diagram commutes:

$$\begin{array}{ccc} \pi_0(\partial\tilde{\mathcal{W}}_n) & \xrightarrow{\tilde{\nu}_\sharp} & \pi_0(\partial\tilde{\mathcal{W}}_n) \\ \downarrow & & \downarrow \\ \pi_0(\partial\mathcal{W}_n^\times) & \xrightarrow{\nu_\sharp^\times} & \pi_0(\partial\mathcal{W}_n^\times) \end{array}$$

Furthermore, suppose that $\tilde{P}_\pm \in \pi_0(\partial\tilde{\mathcal{W}}_n)$ is a pair of oppositely oriented boundary components whose projection $P_\pm^\times \in \pi_0(\partial\mathcal{W}_n^\times)$ satisfies $\nu_\sharp^\times(P_+^\times) = P_-^\times$. Since the covering projections $\tilde{P}_\pm \to P_\pm^\times$ are characteristic modeled on $\tilde{S}_n \to \tilde{S}_{n-1}$, we can promote the gluing free involution $\nu^\times|_{P_\pm^\times}$ to $\tilde{\nu}_{\tilde{P}_\pm}$, namely, the following diagram commutes:

$$\begin{array}{ccc} \tilde{P}_\pm & \xrightarrow{\tilde{\nu}} & \tilde{P}_\mp \\ \downarrow & & \downarrow \\ P_\pm^\times & \xrightarrow{\nu^\times} & P_\mp^\times \end{array}$$

Promoting $\nu^\times$ to all the components of $\partial\tilde{\mathcal{W}}_n$, which have been paired up by $\tilde{\nu}_\sharp$, we obtain an orientation-reversing free involution

$$\partial\tilde{\mathcal{W}}_n \xrightarrow{\tilde{\nu}} \partial\tilde{\mathcal{W}}_n.$$

Finally, we take the claimed finite cover of M_n' to be

$$\tilde{M}_n = \tilde{\mathcal{W}}_n / \tilde{\nu}.$$

The claimed base point $\tilde{o}_n$ of $\tilde{M}_n$ can be chosen as any lift of o_n' in $\tilde{\mathcal{W}}_n(\beta_0)$. It is clear from the construction that the claimed properties of Lemma 5.4 are satisfied by $\tilde{S}_n$ and $(\tilde{M}_n, \tilde{o}_n)$. This completes the induction. □

5.3. **Verification of the asserted properties.** To briefly summarize what we have done so far, under the assumptions of Proposition 5.1, for any given constant $0 < \epsilon < 1$, a sufficiently large positive integer d has been chosen as guaranteed by Lemma 5.2. Moreover, the following commutative diagram of covering maps between base-pointed covers of M has been constructed by Lemma 5.4, where the upper row is the asserted tower of Proposition 5.1 and the lower row is the cyclic tower dual to S encoded by the generalized digital expansion with respect to d:

$$\begin{array}{ccccccccc} \cdots & \longrightarrow & (\tilde{M}_n, \tilde{o}_n) & \longrightarrow & \cdots & \longrightarrow & (\tilde{M}_2, \tilde{o}_2) & \longrightarrow & (\tilde{M}_1, \tilde{o}_1) & \longrightarrow & (M, *) \\ & & \downarrow & & & & \downarrow & & \downarrow & & \downarrow \text{id} \\ \cdots & \longrightarrow & (M_n', o_n') & \longrightarrow & \cdots & \longrightarrow & (M_2', o_2') & \longrightarrow & (M_1', o_1') & \longrightarrow & (M, *) \end{array}$$

As before, we denote by $W_n(\beta)$ the block of pieces of M_n' and $\tilde{\mathcal{W}}_n(\beta)$ their preimage in $\tilde{M}_n$ accordingly.

It remains to verify the requirements of Proposition 5.1 are satisfied. The injectivity radii of $\tilde{M}_n$ at $\tilde{o}_n$ tends to infinity as n grows, because of the second property of Lemma 5.4. The estimate of lifts of ι_Σ is a consequence of the third property of Lemma 5.4. In fact, for every block $\beta \in \mathscr{B}_n$ of level ∞, any component $\tilde{W}$ of $\tilde{\mathcal{W}}_n$ is a regular finite cover of V isomorphic

the pull-back of the characteristic finite cover $\tilde{S}_n \to S$ via the composition $\rho_V : V \to S$ of the retraction ρ of M to S with the inclusion of V into M. Since $\rho \circ \iota_\Sigma$ is homotopic to a constant point map, every elevation of Σ into $\tilde{W}$ is a lift, and the number of such lift is the covering degree $[\tilde{S}_n : S]$. Therefore, the number of lifts of ι_Σ into $\tilde{\mathcal{W}}_n(\beta)$, for any level–$\infty$ blocks $\beta \in \mathscr{B}_n$, equals

$$|\pi_0(\tilde{\mathcal{W}}_n(\beta))| \cdot [\tilde{S}_n : S] \;=\; [\tilde{M}_n : M'_n].$$

The number of lifts of ι_Σ into $\tilde{M}$ is at least the number of such lifts, which is at least

$$(1-\epsilon) \cdot d^{s_n} \cdot [\tilde{M}_n : M'_n] \;=\; (1-\epsilon) \cdot [\tilde{M}_n : M],$$

as asserted. The estimate of the homological torsion size follows immediately from the topological fact:

Lemma 5.5. *If $\Sigma_1, \cdots, \Sigma_m$ are mutually disjointly embedded closed subsurfaces of odd Euler characteristic in an orientable compact 3-manifold N, then $H_1(N;\mathbb{Z})$ contains a submodule isomorphic to $\mathbb{Z}_2^{\oplus m}$.*

Remark 5.6. The author thanks I. Agol for pointing out to him the connection between this fact and exponential torsion growth in dual-graph covers. Theorem 1.3 grows out of that interesting observation.

Proof. Each $[\Sigma_i]$ represents an element $[\Sigma_i] \in H^1(N;\mathbb{Z}_2)$ via Poincaré duality. Since Σ_i has odd Euler characteristic, in the cohomology ring of $\mathbb{Z}_2$ coefficients, $[\Sigma_i]^3$ is nontrivial by [HWZ, Theorem 4.1], but $[\Sigma_i][\Sigma_j]$ is 0 for every distinct pair i, j by the mutual disjointness. It follows that $[\Sigma_i]$ cannot be be lifted to $H^1(N;\mathbb{Z})$ but they span an m-dimensional subspace of $H^1(N;\mathbb{Z}_2)$. By the Universal Coefficient Theorem, $H_1(N;\mathbb{Z})$ contains a torsion submodule isomorphic to $\mathbb{Z}_2^{\oplus m}$. □

This completes the proof of Proposition 5.1.

6. Application to uniform lattices of PSL(2, ℂ)

In this section, we derive the existence of exhausting nested sequence of finite index subgroup with exponential homological torsion growth for uniform lattices of $\mathrm{PSL}(2,\mathbb{C})$, proving Theorem 1.3.

Let Γ be any uniform lattice of $\mathrm{PSL}(2,\mathbb{C})$. Take a torsion-free finite-index subgroup $\dot{\Gamma}$ of Γ, so the quotient space

$$N = \mathbf{H}^3 \,/\, \dot{\Gamma}$$

is an orientable closed hyperbolic 3-manifold.

By [Ag1] and [Wi, Theorem 16.6], we may assume that $\pi_1(N) \cong \dot{\Gamma}$ is cocompactly special, namely, it admits a special cocompact action on a CAT(0) cube complex. By Theorem 1.1, there is a connected closed surface of odd Euler characteristic Σ which admits a π_1–injective quasi-Fuchsian immersion into N. Take S to be an orientable connected closed surface which also admits a π_1–injective quasi-Fuchsian immersion into N, for example, some finite cover of Σ. Form a 2-complex X by attaching an arc α between S and Σ, so

$$\pi_1(X) \cong \pi_1(S) * \pi_1(\Sigma).$$

By well known constructions, we can find a map

$$f : X \longrightarrow N$$

which embeds $\pi_1(X)$ into $\pi_1(N)$ as a quasi-convex subgroup. Moreover, we can require the restrictions to S and Σ to be the claimed quasi-Fuchsian immersions above.

Since $\pi_1(N)$ is word hyperbolic and special, and $\pi_1(X)$ is embedded as a quasi-convex subgroup, it follows from [HW, Theorem 7.3] (see also [Wi, Theorem 4.13]) that $\pi_1(X)$ is a virtual retract of $\pi_1(N)$. In terms of maps, there exists a finite cover

$$M \to N$$

into which f lifts to be a map

$$\iota_X \colon X \longrightarrow M,$$

and there exists a map

$$\rho_X \colon M \longrightarrow X,$$

such that $\rho_X \circ \iota_X$ is homotopic to the identity. In fact, it is not hard to argue that ι_X can be homotoped to be an embedding. Fix a generic base point $*$ of M. Denote by ι_S and ι_Σ the induced embeddings of S and Σ into M, and by ρ the composition of retraction maps:

$$M \xrightarrow{\rho_X} X \xrightarrow{\text{retr.}} S.$$

Finally, we apply Proposition 1.3 to conclude that $(M, *)$ admits an exhausting nested tower of base-pointed finite covers with exponential homological torsion growth. Since $\pi_1(M, *)$ is isomorphic to a finite index subgroup of Γ, the same conclusion holds for Γ as well. This completes the proof of Theorem 1.3.

References

[AbGN] M. Abert, T. Gelander, and N. Nikolov, *Rank, combinatorial cost and homology torsion growth in higher rank lattices*, preprint, 28 pages, 2015. Available at `arXiv:1509.01711v1`.

[Ag1] I. Agol, *The virtual Haken conjecture*, with an appendix by I. Agol, D. Groves, J. Manning, Documenta Math. **18** (2013), 1045–1087.

[Ag2] ———, *Virtual properties of 3-manifolds*, in: Proceedings of the International Congress of Mathematicians, Seoul, Korea, 2014, Vol. 1, pp. 141–170.

[AGM] I. Agol, D. Groves, J. Manning, *An alternate proof of Wise's Malnormal Special Quotient Theorem*, preprint, 42 pages, 2014. Available at `arXiv:1405.0726v2`.

[BV] N. Bergeron and A. Venkatesh, *The asymptotic growth of torsion homology for arithmetic groups*, J. Inst. Math. Jussieu **12** (2013), 391–447.

[DLW] P. Derbez, Y. Liu, S. Wang, *Chern–Simons theory, surface separability, and volumes of 3-manifolds*, J. Topol. (2015), jtv023.

[HW] F. Haglund, D. T. Wise, *Special cube complexes*, Geom. Funct. Anal. **17** (2008), 1551–1620.

[HWZ] C. Hayat-Legrand, S. Wang, H. Zieschang, *Degree-one maps onto lens spaces*, Pacific J. Math. **176** (1996), 19–32.

[Ha] U. Hamenstädt, *Incompressible surfaces in rank one locally symmetric spaces*, Geom. Funct. Anal. **25** (2015), 815–859.

[KM1] J. Kahn and V. Markovic, *Immersing almost geodesic surfaces in a closed hyperbolic 3-manifold*, Ann. Math. **175** (2012), no. 3, 1127–1190.

[KM2] ———, *The good pants homology and a proof of the Ehrenpreis conjecture*, Ann. Math. **182** (2015), no. 1, 1–72. preprint, 2013, 77 pages, `arXiv:1101.1330`.

[KW] K. Kimoto and M. Wakayama, *Equidistribution of holonomy restricted to a homology class about closed geodesics* Forum Math., **14** (2002), no. 3, 383–403.

[LM] Y. Liu and V. Markovic, *Homology of curves and surfaces in closed hyperbolic 3-manifolds*, Duke Math. J. **164** (2015), 2723–2808.

[Lü] W. Lück, *L^2-Invariants: Theory and Applications to Geometry and K-Theory*, Ergebnisse der Mathematik und ihrer Grenzgebiete. 3. Folge. Springer-Verlag, Berlin, 2002.

[PS] M. Pollicott and R. Sharp, *Asymptotic expansions for closed orbits in homology classes*, Geom. Dedicata **87** (2001) no. 1-3, 123–160.
[PW] P. Przytycki and D. T. Wise, *Separability of embedded surfaces in 3-manifolds*, Compositio Mathematica **150** (2014), no. 09, 1623–1630.
[Sa] D. Šarić, *Complex Fenchel-Nielsen coordinates with small imaginary parts* Trans. Amer. Math. Soc Society **366** (2014), 6541–6565.
[SW] P. Sarnak and M. Wakayama, *Equidistribution of holonomy about closed geodesic*, Duke Math. J. **100** (1999), 1–57.
[Su1] H. Sun, *Virtual homological torsion of closed hyperbolic 3-manifolds*, J. Diff. Geom. **100** (2015), 547–583.
[Su2] ———— *Virtual domination of 3-manifolds*, Geom. Topol. **19** (2015), 2277–2328.
[Wi] D. T. Wise, *From Riches to Raags: 3-Manifolds, Right-Angled Artin Groups, and Cubical Geometry*, CBMS Regional Conference Series in Mathematics, vol. 117, Washington, DC, 2012.

Beijing International Center for Mathematical Research, No. 5 Yiheyuan Road, Haidian District, Peking University, Beijing 100871, China P.R.

E-mail address: liuyi@math.pku.edu.cn